

Neural Representation of Space in Rats and Robots

David S. Touretzky, Hank S. Wan, and A. David Redish
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213-3891

Abstract—We describe a computer model that reproduces many observed features of rat navigation behavior, including response properties of place cells and head direction cells. We discuss issues that arise when implementing models of this sort on a mobile robot.

I. Rat Navigation

As they navigate through their environment, rats appear to be employing several types of spatial representations. One type defines “places” based on the views they afford of distal landmarks [14]. Place cells in hippocampus, which fire when the rat is in a particular region of space, are known to be sensitive to visual cues (see [21] for a review). Rats’ sense of place, as reflected in their navigation behavior, has also been shown to rotate in synchrony with the rotation of landmarks, but they fail to recognize the environment when landmarks are permuted [29]. This suggests that the animal’s sense of place is not based on single landmarks but rather on landmark configurations. A number of computer models of visually-driven place cells have been described [3, 26, 27, 28, 33]. Any such model based on purely visual inputs, in which cells are dependent on more than one landmark, should show the effects described in [29].

However, place cells are not driven purely by visual inputs. Most cells still fire correctly when prominent landmarks are removed [22, 23]. As long as some visual cues remain, it is possible that these are sufficient to drive a subset of the place cells which in turn drive the others, “filling in” for the excitation provided by the missing landmarks. Thus, local view-based place descriptions could be reconstructed from partial visual cues. This suggestion is supported by models of the hippocampus as an associative memory [16, 25], and is bolstered by the observation of extensive recurrent connections in CA3. On the other hand, place cell firing is not all-or-none; the firing rate drops off roughly as a gaussian function of distance from the place field’s center. Thus, as the animal travels through the environment, the hippocampal activity pattern changes smoothly, with activity in individual cells increasing and then dropping off again with the approach to and departure from the centers of their fields. The continuous nature of place coding suggests that if there is an associative memory aspect to hippocampal function, the mechanism cannot be as simple as the discrete attractor models studied in the ANN literature [11]. The difficulty of this problem is underscored by the observation that the existing place cell models in the literature (including our own)

are based on essentially feedforward¹ circuitry, while the associative memory models of hippocampus have recurrent connections but do not show place cell-like responses.

Place cells are also known to fire when the animal locomotes in the dark [24]. Thus, visual inputs cannot be the sole means of driving a hippocampal auto-associator. While it is conceivable that other perceptual modalities contribute to sense of place, experiments that control for auditory, olfactory, and tactile cues suggest that they do not play a crucial role in the rat's navigational abilities. In fact, rats can be quite resistant to using olfactory and tactile cues to define goal locations [6].

Hippocampal place coding is primarily a topological representation:² it defines a set of distinct places, augmented by information about which places are adjacent to others. A second spatial representation proposed for rats is a metrical one. Behavioral studies show that rodents traveling in the dark can execute a straight-line path back to their starting location after following a complex, twisty trajectory in the performance of a search task [18]. This is cited as evidence for a path integration faculty in rodents, which other types of animals are also known to possess [10]. Experiments with passive transport show that hamsters' displacement errors when returning home match the distance and direction of the transport [9], which would not be true if they were homing by olfactory or other perceptual cues.

In [8] it was shown that gerbils traveling to a goal location defined by a visual landmark will complete their intended trajectory when the lights are turned off enroute. And in [7] gerbils trained to travel to a goal location indicated by the lit member of an LED array continued on their selected path even when the LED and room lights were switched off. This suggests that the gerbil computes a distance and direction to travel (both metrical quantities), and executes this plan independent of subsequent sensory inputs.

A metrical representation of space presupposes a coordinate system. A polar reference frame might seem the most natural choice, but Gallistel [10] argues that a Cartesian system offers better computational stability when computing incremental position updates, as a path integrator must do.

Metrical representations can be used for trajectory planning, thereby explaining results where a rat traveling directly to a goal is able to navigate through parts of the environment it has not previously visited [12]. Place codes, being purely topological, would require some sort of graph search to do path planning, and would not be able to generate new shortcut

¹There are inhibitory feedback connections in some of these models, and gated recurrent paths in others. However, the computation of place cell activity values does not involve the sort of iterative relaxation and pattern completion found in an auto-associator.

²There is a local metrical structure to place codes in that the firing rates of individual cells fall off roughly as a gaussian function of distance from the center of the place field, and one can estimate the distance between two nearby locations from the overlap of their place code patterns. But bearing information cannot be recovered, and one cannot measure distance between locations more than about a dozen body-lengths apart, due to a lack of place code overlap.

paths through locations the animal had never experienced [32]. Another possible use for metrical representations is to explain the observation that place fields appear to peak roughly 120 msec ahead of the rat, rather than reflecting its current position [20]. We hypothesize that the path integrator predicts the animal's position based on intended movement, and its input to hippocampus provides an anticipatory bias to place cells.

Neurophysiological evidence for a path integration faculty is weaker than for place coding, although there is some data that lesioning the caudate nucleus can disrupt a return from passive transport task [1]. However, a crucial precursor for path integration is an accurate sense of direction, and this has now been well-demonstrated in single-cell recording experiments. Cells in postsubiculum [30, 31] and parietal/retrosplenial association cortex [4, 5] both show responses tuned to head direction, independent of the animal's location in the environment.

II. Robots As Animal Models

Since current robot navigation algorithms are specialized and brittle compared to animal behavior, one might hope to borrow techniques from nature to improve robot performance [2]. This laudable goal is not easily accomplished. There is no reason why merely building place cells into a robot should give better performance than other approaches, such as constructing occupancy grids [19].

However, one problem with current robot algorithms that might be resolved by animal modeling is the selection of appropriate cues. Rats are sensitive to the stability of cues across trials, and they will learn to ignore a cue if experience shows it to be unreliable. Thus, place cells that are sensitive to the position of a white card along the wall of a circular arena can lose their relationship to the card if the animal learns to view the card as unstable [13]. Algorithms for detecting whatever cues are available in an environment and monitoring their stability and usefulness over time could produce more robust robot behavior. In a laboratory environment, for example, the robot should learn to use desks and file cabinets as landmarks while ignoring the positions of chairs and wastebaskets, as these latter objects might be prominent but their positions are unreliable.

Of course, robots that can't distinguish chairs from desks in the first place will have a hard time learning which collection of features constitutes an unstable cue. One reason robots have so much trouble in real-world environments is their limited perceptual abilities, which can be thrown off by simple things such as lighting changes. If a robot cannot produce detailed, reliable perceptual descriptions, the navigation problem becomes far more difficult and we have no reason to expect animal-like performance.

Flexibility in dealing with unexpected situations is another area where animals presently outperform robots. Gerbils searching for a goal in environments that have been made ambiguous by manipulation of landmarks appear to distribute their search effort among plausible alternative locations [8]. In [34] we show how such a distribution can arise from the operation of a place code model. In general, a hippocampal representation that utilizes

a distributed place code based on angles and distances to salient landmarks is likely to be more suitable for coping with environmental changes than, say, an occupancy grid.

A rather different way of relating rats to robots is to view the robot laboratory as a proving ground for theories of animal cognition. Implementation on a real platform forces one to consider issues easily overlooked in computer simulations. In the real world one cannot count on being able to see all the landmarks all the time; many will be occluded. Furthermore, sensor limitations prevent accurate range measurements to distant objects even when they are in the robot's line of sight.

Perception issues come up again when using robots as animal models. In our view, an important future challenge for roboticists interested in "imitating life" will be to reproduce the perceptual capabilities of particular species. Many mobile robots rely heavily on sonar and light stripe sensors, both of which are only effective at close range and cannot provide the distal landmark information required by theories of rat behavior. Color cameras have greater range and far better acuity than sonar; better even than real rats, but their field of view is rather limited compared with the rat's 300° . But the real problem with video cameras is that image processing is computationally expensive. Even something as simple as calculating real-time optic flow requires more processing power than is practical for a mobile robot. Yet optic flow is known to be computed in the early stages of mammalian vision. Such observations underscore the tremendous gulf that remains between today's digital computers and real nervous systems.

III. A Computational Model

In this section we briefly describe our computational model of rat navigation based on coupled mechanisms for place recognition, path integration, and maintenance of head direction, as shown in Figure 1. The point we want to emphasize is that this one architecture reproduces a wide variety of experimental results, as shown in Table 1. Those not covered here are discussed in [34, 35].

Place units in our model compute a fuzzy conjunction of internal and external states. External state information, provided by perception, consists of egocentric angles and distances to visible landmarks. Internal state information comes in the form of coordinates giving the animal's position with respect to one or more reference points. These coordinates are maintained by the path integrator as the animal moves through the environment. A similar proposal was recently made in [17]. Reference points in our model need not be visible objects; they need not even be perceptually distinctive places, although a location such as the animal's nest would be a natural choice. They could simply be locations where something significant once occurred, e.g., the spot where the animal was first released into the experimental chamber, or a place where food was discovered.

Individual place units are tuned to coordinates with respect to one reference point, and angle and distance information with respect to two or more landmarks. The firing rate of the unit reflects the degree of match between its current inputs and the values it is tuned for.

Figure 1: Suggested functional organization of orientation and recognition mechanisms in the rat. 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 to reference point p . (x_g, y_g) are coordinates of goal location relative to some reference point g .

Our units therefore respond as radial basis functions. We use the term “fuzzy conjunction” because terms can drop out when information is not available. So, for example, when the rat navigates in the absence of visual cues, place units respond solely to path integrator coordinates. Conversely, if the animal is dropped into the environment at a random spot, so that path integrator coordinates are not available, the place units are driven solely by perceptual inputs – either angles between pairs of landmarks, or if head direction information is available, allocentric bearings³ to landmarks. Once the animal determines its location it can recall the coordinates associated with that place and use them to reinitialize the path integrator.

Our model uses allocentric bearings to disambiguate visually similar objects, such as the

³Allocentric bearings are measured with respect to the environment, independent of the animal’s orientation. Egocentric bearings are measured relative to the animal’s midline.

Table 1: Features of rat navigation covered by our model.

Experimental Result	Aspect of Model
Place fields controlled by visual landmarks.	Place units are tuned to landmark angles and distances.
Place fields persist in the dark.	Place units are driven by the path integrator.
Place fields are peaked ahead of the rat.	Anticipatory bias from path integrator.
Visually similar features are not confused in un-rotated environments.	Place units can reference allocentric bearings.
Head direction reset by visual cues.	Place units can recall learned allocentric bearings.
Repeated disorientation causes rats to not reset head direction based on visual cues.	Saliency of head direction cues is decreased when retrieved head direction leads to failed predictions.
Place cells can develop direction sensitivity.	Reference point selection based on behavioral context.
Loss of place cell direction sensitivity in response to novelty.	Alerting response activates additional reference points.

identical cylinders used in some of Collett et al.'s landmark-based navigation experiments [8], or the four corners of a rectangular experimental chamber [6, 15]. Head direction information is used to convert egocentric to allocentric directions. Normally head direction is maintained by integrating vestibular cues; when visual input is available, additional help might be derived from optic flow. When the animal is rotationally disoriented so that head direction information is not available, remembered allocentric bearings of landmarks can be used to reconstruct it. Drift in the head direction system resulting from cumulative integration errors can also be corrected this way.

IV. Navigating by Geometric Cues

Cheng [6, 10] showed that rats rely heavily on the geometric structure of their environment to locate a goal, ignoring other types of cues even when these provide more reliable information. He constructed a rectangular experimental chamber with featureless walls and pine chips on the floor, but with a distinctive panel in each corner. Each panel was made of a different material, providing a different texture and visual appearance. Two of the panels had small holes behind which was a cotton wad impregnated with anise in one case and peppermint in the other. The panels also had zero to three pinholes with light shining out of them. In summary, the four panels were as different as possible; to a human they would constitute blatantly obvious place cues.

In one series of experiments, the food reward was always located in front of a particular panel. (Different panels were assigned to different rats.) Cheng rotated the chamber by an arbitrary amount on each trial before introducing the rat, but did not disorient the animal or make any other changes to the apparatus. The result was that rats searched equally often at the correct corner and its rotational equivalent, diagonally opposite it in the chamber,

despite the fact that these two corners had very different panels. His explanation for this result was that the rat focused only on the geometric properties of the rectangular chamber. If the goal was in a corner with a long wall on the left and a short wall on the right, the diagonally opposite corner also satisfied that description. In cases where the animal first went to the incorrect corner, upon failing to find the reward it would often switch over to the diagonally opposite corner, as if realizing that it had made a rotational error.

In a second series of experiments Cheng divided the floor of the chamber into a 9×9 grid forming 81 rectangles, and trained the animal on a working memory task. On each trial, food was buried under the pine chips in one rectangle chosen at random from the set of 80 (the central rectangle was not used.) The rat was then introduced into the chamber and shown where the food was buried. It was allowed to eat some of it before being removed for 75 seconds. During this delay interval the chamber was replaced by an identical copy (to eliminate odor cues), except that the orientation of the new chamber did not match that of the original. The rat was then allowed to search for the food in the new chamber, which should have been indistinguishable from the old except for rotation. Cheng measured the frequency of searching in the correct rectangle (C), the rotationally equivalent rectangle diagonally opposite the correct one (R), or some other spot, which was scored as an error (E). For rotations less than 60° the animals usually chose the correct spot, while for rotations greater than 120° they preferred the diagonally opposite one (Figure 2.) For 90° rotations they chose both spots with equal frequency.

The obvious explanation for this result is that the rat relies on allocentric bearings to disambiguate environmental features that are geometrically indistinguishable. It might, for example, recall the goal location as being at a certain angle and distance from a corner with a long wall on the left and a short wall on the right. In the rotated chamber it must choose one of the two corners meeting this description. The rat fails to make use of nongeometric features of the panels that distinguish the two corners, and instead chooses the one whose allocentric bearing most closely matches the remembered value.

Our model shows that Cheng's results do not require any special cognitive operations on the part of the rat; they can occur as a natural consequence of the place recognition and head direction maintenance systems. When the animal is shown where the food is hidden at the start of a trial, it goes to that location and needs to remember its position so it can return to the goal after the delay period. It does this by noting the path integrator coordinates associated with this place.

Figure 3 shows the response of a place unit whose field is roughly opposite the goal location. This unit is tuned to the bearing and distance of two landmarks (corners), and it has associated with it a set of allocentric coordinates $\langle x_p, y_p \rangle$ measured with respect to its controlling reference point. When the animal is sitting within this unit's place field, the difference between its current path integrator coordinates (which this unit helps estimate) and the remembered coordinates of the goal tell it the path to take to return to the goal.

Now consider Figure 4, showing the place field of the same unit as Figure 3 after the animal is introduced to the second chamber, which has been rotated by a small amount

Figure 2: Cheng's experiment showing percentage of trials in which the rat searched in the **C**orrect rectangle, the **R**otationally equivalent rectangle, or an **E**rror location, as a function of amount of rotation of the experimental chamber. Filled circle indicates the goal. After [6]. Primed values are results from our simulation.




Figure 3: Place field of a unit during in the first part of the working memory task, prior to rotation. The center of this field is in the corner opposite the goal location G.




Figure 4: Place field with a small ($< 60^\circ$) anti-clockwise rotation of the chamber.




Figure 5: Place field with a large ($> 120^\circ$) anti-clockwise rotation of the chamber.

relative to the first. The allocentric bearings of corners no longer closely match the values the place units are tuned for; therefore the response of place units is weaker. The active place units generate predictions about the allocentric bearings of landmarks, and from the discrepancy between predicted and observed bearings the model derives an error signal that it can use to reset its sense of direction. The animal then sees familiar landmarks in the right locations, and can use path integrator coordinates to calculate the direction to the goal.

Figure 5 shows the place field of the same unit in our simulation after a large rotation of the chamber. In this case, each of the four corners is now further from its learned allocentric bearing than the diagonally opposite corner, so the place field flips. Our selected unit's place field is now on the same side of the chamber as the goal. When the animal resets its head direction and calculates the goal location with respect to its assumed position in path integrator coordinates, it will end up searching the rotationally equivalent location.

The middle portion of Figure 2 shows a case where the animal searched the correct location and its rotational equivalent with equal frequency. In our model, place units with fields in the correct and diagonally opposite locations match the corners' allocentric bearings equally poorly. Therefore both sets of units become active. Each generates a different set of allocentric bearing predictions for the observed landmarks, which in turn leads to two distinct values derived for head direction. The animal chooses among competing directions stochastically based on their degree of support. After head direction reset, only one set of place units matches the allocentric bearings well; the losing set becomes inactive. In the maximally ambiguous case of a 90° rotation, the competing head direction values receive nearly identical support, so the animal chooses the two values roughly equally often.

The simulation results reported in Figure 2 show good qualitative agreement with Cheng's data, i.e., for small rotations the correct spot is preferred over the rotational equivalent, while for large rotations the reverse is true. However, Cheng's data is not symmetric: the ratio of correct to incorrect responses is 13:1 for small rotations, while for large rotations the ratio of incorrect to correct responses is only 2:1. This suggests that the rat is probably paying some attention to the panel features, but their influence is not strong enough to consistently override the conflicting information from geometric cues.

Our theory predicts that we will see head direction reset take place if we record from cells in postsubiculum or parietal cortex as the animal is introduced to the rotated chamber. Furthermore, if the animal enters the chamber within a particular cell's place field, that cell should begin firing prior to any shift in head direction cell response. Simultaneous recording from place and head direction cells has been achieved in McNaughton's laboratory [17], so it is feasible to test this prediction in the near future.

Our theory also suggests that we can predict whether the rat will succeed or make an error on each trial by reading the hippocampal place code shortly after it enters the chamber. (McNaughton has managed to accurately predict a rat's position in a familiar environment by simultaneously recording from 80 hippocampal cells using an array of tetrodes [17].) If the place code correctly reflects the animal's location, the proper spot should be searched. If the place fields appear flipped, i.e., the pattern of activity in hippocampus matches that

seen when the animal is in the diagonally opposite location of the unrotated chamber, the animal should make a rotational error.

V. Discussion

We have constructed a general rat navigation architecture based on coupled mechanisms for place recognition, path integration, and maintenance of head direction. A key result of this work is the realization that the architecture that maintains the animal's awareness of position and orientation in space suffices for reproducing results from a wide variety of behavioral experiments, including navigation in the absence of visual input [35], open-field landmark-based navigation tasks [34], and as described in the present paper, navigation in geometrically ambiguous environments. The model is also consistent with single-cell recording data from hippocampal place cells as well as postsubicular and parietal head direction cells.

Our place units are considerably more complex than those of previous hippocampal simulations. They compute fuzzy conjunctions of internal and external state variables, with terms temporarily dropping out when information is unavailable. The units recognize places using any combination of allocentric bearings to landmarks (when head direction information is available), distances to landmarks, egocentric angular difference between pairs of landmarks, and coordinates with respect to the unit's controlling reference point as estimated by the path integrator. The units learn the allocentric bearings of landmarks visible within their place fields and use this information to guide place recognition. It could also be used to correct for drift in the head direction system.

The sophistication of these place units, crucial to the behavioral richness of the model, suggests that they should not be identified with single pyramidal cells. The hippocampal complex contains several types of neurons, and its connections with entorhinal cortex and subiculum form a complex information processing loop. Thus, we see our place units as representing small bits of circuitry in which CA3 and CA1 pyramidal cells (place cells) are but one component.

Acknowledgements

We thank Bruce McNaughton and Neil Burgess for helpful discussions. Hank Wan and David Redish were supported by NSF fellowships.

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] I. A. Bachevalier and A. M. Waxman. Mobile robot visual mapping and localization: a view-based neurocomputational architecture that emulates hippocampal place learning. *Neural Networks*, 7, 1994. in press.
- [3] 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.
- [4] L. L. Chen. *Head-directional information processing in the the rat posterior cortical areas*. PhD thesis, University of Colorado, 1991.
- [5] L. L. Chen, B. L. McNaughton, C. A. Barnes, and E. R. Ortiz. Head-directional and behavioral correlates of posterior cingulate and medial prestriate cortex neurons in freely-moving rats. *Society of Neuroscience Abstracts*, 16:441, 1990.
- [6] K. Cheng. A purely geometric module in the rat’s spatial representation. *Cognition*, 23:149–178, 1986.
- [7] T. S. Collett. The use of visual landmarks by gerbils: Reaching a goal when landmarks are displaced. *Journal of Comparative Physiology A*, 160:109–113, 1987.
- [8] 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.
- [9] 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.
- [10] C. R. Gallistel. *The Organization of Learning*. MIT Press, Cambridge, MA, 1990.
- [11] J. Hertz, A. Krogh, and R. G. Palmer. *Introduction to the Theory of Neural Computation*. Addison-Wesley, Reading, MA, 1991.
- [12] J. R. Keith and K. M. McVety. Latent place learning in a novel environment and the influences of prior training in rats. *Psychobiology*, 16(2):146–151, 1988.
- [13] J. J. Knierim, B. L. McNaughton, C. Duffield, and J. Bliss. On the binding of hippocampal place fields to the inertial orientation system. *Society of Neuroscience Abstracts*, 19:795, 1993.
- [14] B. J. Kuipers and T. S. Levitt. Navigation and mapping in large-scale space. *AI Magazine*, 9(2):25–43, 1988.
- [15] J. Margules and C. R. Gallistel. Heading in the rat: Determination by environmental shape. *Animal Learning and Behavior*, 16(4):404–410, 1988.
- [16] B. L. McNaughton. Neuronal mechanisms for spatial computation and information storage. In L. Nadel, L. Cooper, P. Culicover, and R. M. Harnish, editors, *Neural Connections, Mental Computation*, chapter 9, pages 285–350. MIT Press, Cambridge, MA, 1989.
- [17] 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 Neurosciences*. MIT Press, Boston, 1993.
- [18] M. L. Mittelstaedt and H. Mittelstaedt. Homing by path integration in a mammal. *Naturwissenschaften*, 67:566–567, 1980.

- [19] H. P. Moravec. Sensor fusion in certain grids for mobile robots. *AI Magazine*, 9(2):61–74, 1988.
- [20] R. U. Muller and J. L. Kubie. The firing of hippocampal place cells predicts the future position of freely moving rats. *Journal of Neuroscience*, 9(12):4101–4110, 1989.
- [21] 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.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] E. T. Rolls. The representation and storage of information in neuronal networks in the primate cerebral cortex and hippocampus. In R. Durbin, C. Miall, and G. Mitchison, editors, *The Computing Neuron*. Addison-Wesley, Reading, MA, 1989.
- [26] 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.
- [27] 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.
- [28] P. E. Sharp. Computer simulation of hippocampal place cells. *Psychobiology*, 19(2):103–115, 1991.
- [29] S. Suzuki, G. Augerinos, and A. H. Black. Stimulus control of spatial behavior on the eight-arm maze in rats. *Learning and Motivation*, 11:1–18, 1980.
- [30] 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.
- [31] 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.
- [32] E. C. Tolman. Cognitive maps in rats and men. *Psychological Review*, 55:189–208, 1948.
- [33] A. Treves, O. Miglino, and D. Parisi. Rats, nets, maps, and the emergence of place cells. *Psychobiology*, 20(1):1–8, 1992.
- [34] H. S. Wan, D. S. Touretzky, and A. D. Redish. Computing goal locations from place codes. *Manuscript*, 1994.
- [35] 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, 1994.