

Appeared in *Neural Computation* 5(6):869-884.
This is the Tech Report version, released as Tech Report CMU-CS-93-117.

Neural Representation of Space Using Sinusoidal Arrays

David S. Touretzky¹ A. David Redish Hank S. Wan
March 1993
CMU-CS-93-117

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

To appear in *Neural Computation*.

Abstract

O'Keefe (1991) has proposed that spatial information in rats might be represented as phasors: phase and amplitude of a sine wave encoding angle and distance to a landmark. We describe computer simulations showing that operations on phasors can be efficiently realized by arrays of spiking neurons that re-code the temporal dimension of the sine wave spatially. Some cells in motor and parietal cortex exhibit response properties compatible with this proposal.

¹Address all correspondence to the first author. Electronic mail address: dst@cs.cmu.edu.

This work was supported by a contract from Fujitsu Corporation. Hank Wan and David Redish were supported by NSF Graduate Fellowships.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Fujitsu Corporation, the National Science Foundation, or the U.S. government.

Keywords: Neural Modelling, Spatial Reasoning, Parietal Cortex, Sinusoidal Arrays

1. Introduction

Any vector in polar coordinates $\vec{v} = (r, \phi)$ can be represented as a sine wave $f(t) = r \cos(\omega t + \phi)$, where r is amplitude, ϕ is phase, and ω is (constant) frequency. This is commonly known as a *phasor*. The advantage of phasor representation is that translation and rotation of a vector are both trivial operations. Translation is achieved by addition of sine waves, and rotation can be obtained by phase shifting or temporal delay. O’Keefe (1991) has suggested that rats might use phasors to encode angle and distance to landmarks. In his proposal, hippocampal theta provides the reference signal for determining phase.

This temporal approach to encoding a sine wave has some drawbacks. The 7-12 Hz theta rhythm may be too slow to support real-time spatial reasoning tasks requiring rapid manipulation of phasors. Furthermore, maintaining even a modest angular resolution of 10° relative to a roughly 10 Hz reference signal requires a temporal resolution of 3 msec. Although some specialized sensory systems are known to make much finer discriminations (e.g., acoustic imaging in bats and dolphins, or auditory localization in barn owls), we are reluctant to require this degree of temporal precision at the higher cognitive level associated with spatial reasoning. Instead, we suggest that phasor operations are more plausibly realized by re-coding the temporal dimension of the sine wave spatially, using populations of spiking neurons. We propose an architecture called the *sinusoidal array* for manipulating vectors in phasor form, and report the results of computer simulations.

There is some experimental evidence that sinusoidal array representations may exist in rat parietal cortex and in rhesus motor or parietal cortex. We propose an experiment to test this hypothesis in rats.

2. Sinusoidal Arrays

To encode a phasor as a sinusoidal array, we replace the continuous temporal signal $f(t)$ by a distributed pattern of activity over an array of N elements, as in Figure 1. The value encoded by the i th array element is the amplitude of the sine wave sampled at point $2\pi i/N$. That is, the activity level of the i th array element encoding the vector (r, ϕ) is given by $f(r, \phi, i) = r \cos(\phi + 2\pi i/N)$, for $0 \leq i < N$. Note that for the special case of $N = 4$, the sinusoidal array encoding is exactly the Cartesian encoding $(x, y, -x, -y)$, where $x = r \cos \phi$ and $y = r \sin \phi$.

Each sinusoidal array element is a collection of neurons. Its activity level is encoded by the neurons’ average firing rate, or equivalently, the average percentage of neurons firing at any instant. If the neuronal population is sufficiently large, this representation can encode values with high precision even when individual neurons are noisy and have a limited number of discriminable firing rates.

In order to be able to represent the negative half of the sine wave, neurons in a sinusoidal array fire at a rate $F(r, \phi, i) = k \cdot f(r, \phi, i) + b$, where k is a gain parameter and b the baseline firing rate. In our simulations, the baseline firing rate is 40 spikes/second. This gives the neuron a dynamic range of 0-80 Hz, which is compatible with cells in parietal cortex.

A significant advantage of the sinusoidal array representation is that it allows coordinate transforms to be done nearly instantaneously. If the signal $f(t)$ were represented temporally, the simplest way to determine its phase would be to wait for the peak. But one might have to wait up to one full

Figure 1: The phasor (r, ϕ) and its sinusoidal array representation.

period, which would be 140 msec for a 7 Hz signal. Alternatively, the phase could be determined by estimating the slope $f'(t)$ and taking the arc cosine, but this solution seems less neurally plausible than the spatial encoding investigated here.

A drawback of the sinusoidal array representation is that angular resolution is limited to $2\pi/N$. But even modest values of N appear to give sufficient resolution for navigation tasks. In our simulations we chose $N = 24$, giving an angular resolution of $\pm 7.5^\circ$.

3. Vector Operations with Sinusoidal Arrays

In order to successfully complete a landmark-based navigation task, an animal must perform some coordinate transformations to derive a goal location from observed landmark positions. These transformations include at least translation, probably negation, and perhaps also rotation.

In a phasor-based coordinate system, translation of a vector is accomplished by adding the corresponding sine waves, e.g., $f(t) = f_1(t) + f_2(t)$. In the sinusoidal array representation, translation is accomplished by element-wise linear addition of firing rates: $F(i) = F_1(i) + F_2(i) - b$, for $0 \leq i < N$. The subtraction of one baseline value b normalizes the result; it can be accomplished by giving the summation neuron an inhibitory bias term equal to the baseline firing rate.

Negation of a vector can be accomplished in a variety of ways. Given a maximum activity level M for an element, we can compute $F(i) = M - F_1(i)$ for $0 \leq i < N$ by using inhibitory connections to units whose baseline firing rate is M . However, since negation of a vector is usually required only as part of a vector subtraction operation, the easiest solution may be to use the addition mechanism described in the previous paragraph, but with one of the vectors rotated by 180° . This gives $F(i) = F_1(i) + F_2(i + N/2 \bmod N) - b$.

If translation and negation were the only required operations, there would be no advantage to using phasors. All computations could be done in Cartesian coordinates, using any neural encoding

Figure 2: Schematic diagram of the shifter circuit. The signal entering at right goes through a contrast enhancement stage and winner-take-all phase detector, which determines the amount by which the input sine wave (top left) should be shifted. Light-colored lines indicate lateral inhibition connections. Only a subset of the permutation channel connections are shown.

The 1-of- N phase detector contains one neuron for each of the N sinusoidal array elements. These neurons integrate their inputs over a brief time interval; the one receiving the largest input reaches threshold and fires first. We think of these phase detector neurons as similar to fast-spike inhibitory interneurons. They have small refractory times and two sorts of postsynaptic effect: a short timescale inhibition of other phase detector cells (whose recovery from inhibition initiates a new winner-take-all round), and a long timescale inhibition that acts as a gating signal for channel-

inhibitory neurons in the second half of the shifter.

The shifter’s N gated permutation channels each copy the activity of the N -element input array to the N -element output array, permuting the elements along the way. When the j th channel is active, it copies the activation of input element i to output element $i - j \bmod N$, for $0 \leq i < N$. The channels have associated with them tonically active channel-inhibitory neurons that keep them silent most of the time. These are the same type of inhibitory units as the phase detector neurons, except that their only inputs are inhibitory. When a neuron in the phase detector fires, it inhibits the corresponding channel-inhibitory neuron, thereby dis-inhibiting the channel and allowing the shifted sine wave to appear across the output array.

Anderson and Van Essen describe a shifter using $\log_2 N$ levels where each level has two permutation channels, giving $O(N)$ connections. In a refinement of this model, Olshausen et al. use four levels with varying numbers of nodes, and fan-ins of approximately 1000, mirroring the connectivity of cortical areas V1, V2, V4, and IT. Because our own N is so small ($N = 24$ in the simulations), we can use a single level with N channels and $O(N^2)$ connections. Aside from the obvious advantage of simplicity of connection structure, this allows us to use a simple 1-of- N representation for the amount by which to shift, rather than the more complex binary encoding required by Anderson and Van Essen, or the distributed encoding of Olshausen et al. Our model is also simpler because it requires only shunting inhibition, whereas theirs requires multiplicative synapses.

The shifter circuit is not central to our theory. As discussed in the next section, many rodent navigation tasks can be performed using just translation. However, in situations where the reference frame must be determined anew on each trial based on the orientation of a cue array, there does appear to be a need for mental rotation of some sort. The shifter offers a solution to the general problem of rotation of vectors. But for some navigation tasks, the animal could instead slew its internal compass.

4. Rodent Navigation

In a remarkable series of experiments, Collett, Cartwright, and Smith investigated landmark learning behavior in gerbils [Collett et al. 86]. We will describe two of their simpler experiments here. Figure 3 shows the result of training a gerbil to find a food reward at a constant distance (50 cm) and compass bearing from a cylindrical landmark. The landmark was moved to a different location at each trial. Once trained (requiring on the order of 150 trials), the gerbil proceeded directly to the goal location from any starting location, and spent most of its time searching in the goal location.

To model this behavior, we assume that the gerbil has learned that a constant memory vector M describes the remembered angle and distance of the landmark from the goal. On each trial, the gerbil’s perceptual apparatus produces a vector P that describes the location of the landmark relative to the animal’s current position. Thus, the position of the goal relative to the animal can be computed by vector subtraction: $G = P - M$. Collett et al. show that the animal must be computing this location, rather than simply moving to make its view of the landmark match a stored memory of the goal, by turning off the lights after it had begun moving toward the goal. The animal still proceeded directly to the goal.

The calculation of the goal location relies on a critical assumption: that the memory vector

Figure 3: Learning to find food at a constant distance and bearing from a landmark, after [Collett et al. 86] and [Leonard & McNaughton 90]. S marks sample starting locations; F is food location; solid circle is the landmark; concentric circles show distribution of search time when trying to locate food reward. The majority of search time is spent at the goal location.

M , the perceptual vector P , and the goal vector G share the same reference direction, which we call *global north*. In [Collett et al. 86] this commonality was attributed to “unspecified directional cues.” Recently it has been shown that rodents possess a highly accurate internal compass which allows them to judge head direction even in the absence of visual cues [Chen et al. 90, Chen 91, Taube et al. 90a, Taube et al. 90b]. The compass is not related to magnetic north; it is maintained in part by integrating angular accelerations over time [Mittelstaedt & Mittelstaedt 80, Etienne 87, Markus et al. 90].

McNaughton (personal communication) has observed that rats taken from their home cages will maintain their compass as they are carried into another room for experiments, so they have a consistent reference frame available even if the experimental environment is poor in directional cues.¹ This is significant because an internal compass that provides a stable north across all learning trials allows many simple navigation tasks to be performed without resorting to mental rotation.

In our simulation of the Collett et al. task, we assume that the perceptual and memory systems orient their respective vectors, P and M , using the same internal compass. The sinusoidal array then computes G by vector subtraction.

Figure 4 shows a more demanding experiment in which the cue array is rotated as well as translated on each trial. The bearing of the food reward is constant with respect to the line joining landmarks L_1 and L_2 , as shown in Figure 5, but not with respect to the more salient cue provided by the internal compass. Tasks of this sort, in which bearings must be measured with respect to the cue array, should be more difficult to learn (Collett et al.’s observations support this), and would seem to require mental rotation.

¹However, the compass can be confused if the box used to transport the animal is spun through several revolutions.

Figure 5: Bearing α to food reward is measured with respect to the line joining landmarks L_1 and L_2 , not global north.

Figure 6: Computations involved in solving the task shown in Figure 4. P_1, P_2 : coordinates of landmarks in the perceptual reference frame; M_1, M_2 : remembered coordinates of landmarks as seen from the food location; R : rotational alignment factor; G : computed goal location in the perceptual reference frame.

5. Details of the Computer Simulations

Our computer simulations are based on abstract neuron models that are simpler than compartmental models, but retain many important properties of real neurons, such as spiking behavior.

The simulation uses two types of neurons: pyramidal cells and inhibitory interneurons. Sinusoidal arrays contain 24 elements with 100 pyramidal cells each. Our abstract pyramidal cell has a resting potential of 0, a threshold $\theta = 1$, and a typical fan-in of 20 (but as high as 240 in the shifter) with synaptic weights of 0.1. It sums its inputs over time, and when it reaches threshold, it fires a spike. Spiking is treated as an instantaneous event, i.e., it lasts for one clock tick, after

which the cell zeros its net activation and enters a refractory state. For the experiments reported here, a clock tick, Δt , is 0.1 msec. The cell’s refractory period is 1/80 second, limiting the peak firing rate to 80 Hz. It’s important that the clock rate be significantly faster than peak firing rate, so that inputs are not lost when a cell zeros its net activation. Only impulses arriving at the exact moment a cell spikes will be lost; during the refractory period the cell continues to integrate its inputs.

Pyramidal cells make up the summation module used for addition and subtraction of phasors. Cells in the summation module receive excitatory inputs from two sinusoidal arrays, following the equation $F(i) = F_1(i) + F_2(i) - b$. A neuron in the i th array element will receive inputs from 10 randomly-chosen neurons from the i th element of each input array. The bias term $b = 40$ Hz is implemented by decrementing the net activation by $-b \cdot \theta \cdot \Delta t$ every clock tick, but the total activation of the cell is not permitted to go below 0.

Pyramidal cells also make up the output array of the shifter module. These cells have a fan-in of 240 since they receive 10 inputs from each of N permutation channels. They do not require a bias term.

The second type of model neuron is a fast-spike inhibitory interneuron used in the shifter. Both the phase detector neurons and the permutation channel inhibitory neurons are of this type. It has a resting level of 0 and a threshold of 1, like the pyramidal cell, but the refractory period is only 5 msec. The firing of a phase detector neuron has two distinct effects. First, it inhibits all the other phase detector neurons, essentially setting their net activation to zero. Second, it inhibits the corresponding channel-inhibitory neuron, allowing the permutation channel to open. Lateral inhibition of phase detector cells should have a short time course, so that when a neuron loses the race to fire first it can re-enter the competition in short order. But channel-inhibitory neurons should be inhibited for a relatively long time, because we do not want the channel to close again between successive firings of its controlling phase detector. In cortex, GABA_A inhibition has a short time course, while GABA_B inhibition has a long time course. It therefore does not seem unreasonable to posit different inhibitory effects arising from the same interneuron.

The channel-inhibitory neurons, when not themselves inhibited, shut down the permutation channel. This could be accomplished in real neural systems in several ways. If we assume that the i th channel’s bundle of connections from input cells to a cell in the shifter’s output array are distributed throughout the output cell’s dendritic tree, then shutting down the channel would require inhibitory axo-axonic synapses at many select sites. But if connections comprising the i th channel were localized to a specific region of the output cell’s dendritic tree, the channel-inhibitory interneuron would require only a single synapse onto the base of this subtree. Because our simulation is not at the compartmental level, we do not distinguish between these possibilities in our model.

We add noise to the model by perturbing each cell’s activation level by a small random value at each clock tick. For 5% noise, we use perturbations in the range $\pm 0.025M \cdot \theta \cdot \Delta t$, where M is the cell’s maximum firing rate. Small amounts of noise actually improve the shifter’s performance by preventing the output cells within an array element from synchronizing with each other due to inhibition from contrast enhancement, described below.² Noise also prevents a phase detector cell from consistently winning the race to inhibit its neighbors just because the cells that synapse onto

²Synchronization would cause anomalous behavior in any phase detector that used this signal as input, unless the phase detector cells integrated their inputs over a much longer time period.

it happened to start out with a slightly higher initial activation level.

A technical problem with the shifter suggests that we may want to add basket cells to our model. We found that for the shifter to work correctly, the phase detector must be producing a stable output, i.e., reporting a consistent phase. However, when the sine wave input to the phase detector is of small amplitude, the peak can be difficult to determine precisely, so the phase detector’s output wanders among several nearby values. This results in the opening of different permutation channels at different times, degrading the shifter’s output representation. To prevent this, we introduced a contrast enhancement layer with a form of “center-surround” feedback inhibition to pre-process the phase detector’s input and make the peak easier to find. In real neural systems, this type of inhibitory feedback is thought to be provided by basket cells [McNaughton & Nadel 90]. The details of our model’s contrast enhancement mechanism are a bit ad hoc at present and are in need of refinement, but preliminary results show that it does result in correct and stable phase detector output.

If the inhibitory feedback is set at a high level, the contrast enhancement process yields an array representation with only one active element, thereby anticipating the winner-take-all function of the phase detector. However, for a range of lower values, instead of winner-take-all behavior contrast enhancement produces cells with triangular response functions. The firing rates of these cells peak at a certain preferred direction, fall off roughly linearly within 30-60 degrees of that direction, and are elsewhere flat and close to zero. As discussed in the next section, cells with this behavior have been found in postsubiculum by Taube et al.

We have also run simulations varying the number of neurons in a sinusoidal array. There was no appreciable advantage to doubling the number to 200 neurons per element. There was a slight penalty for using only 50 neurons: it took longer for the shifter to settle down and produce a consistent output signal, because contrast enhancement had to be done more slowly to avoid errors. With 20 neurons per element the system was unstable.

6. Experimental Evidence for Sinusoidal Arrays

A necessary condition for sinusoidal arrays to exist in cortex is the presence of cells whose response pattern obeys the function $F(r, \phi) = b + k \cdot r \cos \phi$, where distance r and angle ϕ are measured either egocentrically or allocentrically.

Georgopoulos et al. have formulated a similar equation, $d(M) = b + k \cos(\theta_{CM})$, to describe the behavior of neurons in rhesus parietal cortex [Kalaska et al. 83] and motor cortex [Georgopoulos et al. 86]. These neurons have firing rates proportional to the cosine of the angle between a “preferred direction vector” C and an intended reaching vector M . Different cells have different directional preferences³ and hence different firing rates for a given movement. Their collective activity forms a “neural population vector” that can express any angle of intended motion.

Another important piece of evidence in support of the sinusoidal array hypothesis is the finding in rats of cells that encode head direction with respect to either a visual landmark or an inertial reference frame. These cells appear to be part of the animal’s internal compass referred to earlier. Taube et al. report head-direction sensitive cells in postsubiculum with sharp directional preferences

³The preferred direction C plays the role of the array position i in our formula for $F(r, \phi, i)$.

Figure 7: Tuning curves for a cell in parietal area Oc2M when the animal is motionless or making a left or right turn. Modified from [Chen 91, p. 118].

The crucial question for both the rat and primate data is whether cells with a sinusoidal response function are also sensitive to distance. Schwartz and Georgopoulos have found this to be the case in rhesus motor cortex [Schwartz & Georgopoulos 87]. They first varied the angle of a constant-distance target in a reaching task, in order to determine the preferred direction for each cell. Subsequently, they varied the distance between the animal and the target when the target was located at the cell's preferred direction. They report a substantial number of direction-sensitive cells with weak but statistically significant linear response as a function of target distance.

In the case of the rat parietal recordings, in order to measure sensitivity to distance the animal would have to be attending to some known location. One way to accomplish this would be to train the rat to perform a landmark-based navigation task as in Figure 3, and then look for direction-sensitive parietal cells whose response varied linearly with distance to either the landmark or the goal.

7. Discussion

Hippocampal theta may play some role as a reference signal for navigation, but it is probably not related to compass direction. O’Keefe and Recce (1992) report that the phase at which place cells fire relative to the theta rhythm varies through 360° as the animal enters, proceeds through, and exits the cell’s place field. This has led Burgess, O’Keefe, and Recce to propose a navigation model in which phase information is used to distinguish entering vs. exiting. In conjunction with head direction information and a separate layer of goal cells, the net firing field of subicular place cells at phase 0° is peaked *ahead* of the rat, allowing the animal to navigate by homing to a goal location [Burgess et al. 93].

The Burgess et al. model has a number of interesting properties, but it cannot deal with complex navigation tasks of the sort Collett et al. have studied, with cue arrays that change position and orientation from trial to trial. While the hippocampus is known to play an important role in spatial behavior, researchers such as [Nadel 91] claim that its role is spatial memory, not planning and navigation. Parietal cortex appears to be involved in these latter tasks [Stein 91].

McNaughton et al. (1991) propose a model of directional sense based on both vestibular sensations and visual cues. In darkness or unfamiliar environments, the animal maintains its compass by inertial means, using an associative memory “table lookup” scheme to compute its new heading from the old heading plus angular acceleration. But in familiar environments, “local view” cells (possibly hippocampal place cells) adjust the compass to agree with the learned heading associated with that view direction.

We mentioned previously that compass slewing might replace the second rotation when performing Collett et al.’s rotating cue array task. McNaughton (personal communication) has suggested that if local view cells can determine compass direction by direct matching of visual landmarks, the first subtraction and rotation steps might also be eliminated, leaving just one vector subtraction. We agree with the notion that distant landmarks should control the animal’s compass in familiar environments. But it seems less plausible that viewing a configuration of *nearby* landmarks would provide sufficiently accurate heading information to solve the rotating cue array task by table lookup, because the view could change significantly with relatively small translations. Hence we believe at least one mental rotation step is required.

Elsewhere in their paper McNaughton et al. speculate that trajectory computations (vector subtractions) might be done by the same table lookup mechanism as they propose for updating the inertial compass. The drawback of this proposal is the large table that would be required to represent all possible pairs of vectors, and the cost of filling in the entries. The sinusoidal array appears to offer a simpler solution for vector arithmetic.

The neural architecture we’ve described is compatible with properties of parietal cortex. It manipulates phasors as sinusoidal arrays, but it does not explain how such representations arise in the first place. We simply assume that perceptual and memory vectors are available in the required form. We defend this assumption by noting that sinusoidal encodings of angles have already been found in rats and monkeys. Indications of a linear sensitivity to distance in rhesus sinusoidal cells reported by Schwartz & Georgopoulos offer additional support. At this point, the most important test of our model is whether rat parietal cells can be found with cosine response functions that are also linearly sensitive to distance.

Two other properties of our model are worth noting. As presently formulated, all cells in a sinusoidal array element have the same preferred direction (as do cells in a single orientation column in visual cortex), so there are only N directions represented. If the preferred directions of real parietal cells are found to cluster into a small number of discrete, evenly-spaced values, this would be strong evidence in favor of the sinusoidal array hypothesis. However, we expect our model would also function correctly using input units with preferred directions smoothly distributed around the circle, so that neurons in bin i had a preferred direction somewhere within $2\pi(i \pm 0.5)/N$. We have not yet verified this experimentally, however. Due to the many-to-one connectivity of pyramidal cells, units in the output sinusoidal array should still show preferred direction values close to the centers of their respective bins.

The model also assigns the same scale factor k to all neurons in an array. But experimenters report a wide range of peak firing rates for direction-sensitive cells in both postsubiculum and parietal cortex. We again expect the model to function correctly under this condition, assuming only that the mean scale factor is the same across elements.

Acknowledgements

This work was supported by a grant from Fujitsu Corporation. Hank Wan and David Redish were supported by NSF Graduate Fellowships. We thank Bruce McNaughton and an anonymous referee for helpful comments on an earlier draft of this paper, and Longtang Chen for permission to reproduce one of his figures.

References

- [Anderson & Van Essen 87] Anderson, C. H., and Van Essen, D. C. (1987) Shifter circuits: a computational strategy for dynamic aspects of visual processing. *Proceedings of the National Academy of Sciences, USA*, **84**, 1148-1167.
- [Burgess et al. 93] Burgess, N., O'Keefe, J., and Recce, M. (1993) Using hippocampal 'place cells' for navigation, exploiting phase coding. In S. Hanson, J. Cowan, and L. Giles (eds.), *Advances in Neural Information Processing Systems 5*. San Mateo, CA: Morgan Kaufmann.
- [Chen et al. 90] Chen, L. L., McNaughton, B. L., Barnes, C. A., and Ortiz, E. R. (1990) Head-directional and behavioral correlates of posterior cingulate and medial prefrontal cortex neurons in freely-moving rats. *Society of Neuroscience Abstracts*, **16**, 441.
- [Chen 91] Chen, L. L. (1991) Head-directional information processing in the rat posterior cortical areas. Doctoral dissertation, University of Colorado.
- [Collett et al. 86] 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.
- [Etienne 87] Etienne, A. S. (1987) The control of short-distance homing in the golden hamster. In P. Ellen & C. Thinus-Blanc (eds.), *Cognitive Processes and Spatial Orientation in Animals and Man*, 233-251. Dordrecht: Martinus Nijhoff.
- [Georgopoulos et al. 86] Georgopoulos, A. P., Schwartz, A. B., and Kettner, R. E. (1986) Neuronal population coding of movement direction. *Science*, **233**, 1416-1419.

- [Kalaska et al. 83] Kalaska, J. F., Caminiti, R., and Georgopoulos, A. P. (1983) Cortical mechanisms related to the direction of two-dimensional arm movements: relations in parietal area 5 and comparison with motor cortex. *Exp. Brain Res.* **51**:247-260.
- [Leonard & McNaughton 90] Leonard, B., and McNaughton, B. L. (1990) Spatial representation in the rat: conceptual, behavioral, and neurophysiological perspectives. In R. P. Kesner and D. S. Olton (eds.), *Neurobiology of Comparative Cognition*, 363-422. Hillsdale, NJ: Erlbaum.
- [Markus et al. 90] Markus, E. J., McNaughton, B. L., Barnes, C. A., Green, J. C., and Meltzer, J. (1990) Head direction cells in the dorsal presubiculum integrate both visual and angular velocity information. *Society for Neuroscience Abstracts* **16**, 441.
- [McNaughton & Nadel 90] McNaughton, B. L., and Nadel, L. (1990) Hebb-Marr networks and the neurobiological representation of action in space. In M. A. Gluck and D. E. Rumelhart (eds.), *Neuroscience and Connectionist Theory*, 1-63. Hillsdale, NJ: Erlbaum.
- [McNaughton et al. 91] McNaughton, B. L., Chen, L. L., and Markus, E. J. (1991) "Dead reckoning," landmark learning, and the sense of direction: a neurophysiological and computational hypothesis. *Journal of Cognitive Neuroscience* **3**(2):190-202.
- [Mittelstaedt & Mittelstaedt 80] Mittelstaedt, M.-L., and Mittelstaedt, H. (1980) Homing by path integration in a mammal. *Naturwissenschaften* **67**:566-567.
- [Nadel 91] Nadel, L. (1991) The Hippocampus and space revisited. *Hippocampus* **1**(3):221-229.
- [O'Keefe 91] O'Keefe, J. (1991) An allocentric spatial model for the hippocampal cognitive map. *Hippocampus* **1**(3):230-235.
- [Olshausen et al. 92] Olshausen, B., Anderson, C., and Van Essen, D. (1992) A neural model of visual attention and pattern recognition. CNS Memo 18, Computation and Neural Systems Program, California Institute of Technology.
- [Schwartz & Georgopoulos 87] Schwartz, A. B., and Georgopoulos, A. P. (1987) Relations between the amplitude of 2-dimensional arm movements and single cell discharge in primate motor cortex. *Soc. Neurosci. Abstr.* **13**:244.
- [Stein 91] Stein, J. (1991) Space and the parietal association areas. In J. Paillard (ed.), *Brain and Space*, pp. 185-222, New York: Oxford University Press.
- [Taube et al. 90a] Taube, J. S., Muller, R. I., and Ranck, J. B. J. (1990a) Head direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, **10**, 420-435.
- [Taube et al. 90b] Taube, J. S., Muller, R. I., and Ranck, J. B. J. (1990a) Head direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *Journal of Neuroscience*, **10**, 436-447.