Separating hippocampal maps

A. David Redish* and **David S. Touretzky**†
Computer Science Department and
Center for the Neural Basis of Cognition
Carnegie Mellon University
Pittsburgh, PA 15213-3891

October 21, 1997

Abstract

The place fields of hippocampal cells in old animals sometimes change when an animal is removed from and then returned to an environment [Barnes *et al.*, 1997]. The ensemble correlation between two sequential visits to the same environment shows a strong bimodality for old animals (near 0, indicative of remapping, and greater than 0.7, indicative of a similar representation between experiences), but a strong unimodality for young animals (greater than 0.7, indicative of a similar representation between experiences).

One explanation for this is the *multi-map* hypothesis in which multiple maps are encoded in the hippocampus: old animals may sometimes be returning to the wrong map. A theory proposed by Samsonovich and McNaughton (1997) suggests that the Barnes *et al.* experiment implies that the maps are pre-wired in the CA3 region of hippocampus. Here, we offer an alternative explanation in which orthogonalization properties in the dentate gyrus (DG) region of hippocampus interact with errors in self-localization (reset of the path integrator on re-entry into the environment) to produce the bimodality.

To appear in: Spatial Functions of the Hippocampal Formation and the Parietal Cortex, edited by Neil Burgess, Kathryn Jeffery, and John O'Keefe. Oxford University Press.

*Author's current address: NSMA, University of Arizona, Life Sciences North Bldg, Rm. 384, PO Box 24-5115, Tucson AZ 85724. Electronic mail address: adr@nsma.arizona.edu
†Electronic mail address: dst@cs.cmu.edu

1 Introduction: Reference frames.

Hippocampal place cells also show correlations to non-spatial aspects of the world, including environment [Kubie and Ranck, 1983, Thompson and Best, 1989, Muller and Kubie, 1987], task within environment [Markus *et al.*, 1995], and even stage within task [Eichenbaum *et al.*, 1987, Eichenbaum and Cohen, 1988, Cohen and Eichenbaum, 1993, Hampson *et al.*, 1993, Gothard *et al.*, 1996]. Some have argued that these experiments imply that place cells should be understood as being general context cells, with space being only one of many parameters to which they are tuned [Eichenbaum *et al.*, 1992, Wiener, 1993, Eichenbaum, 1996].

However, even when a cell shows different firing patterns under two conditions, the cell still show place fields under both conditions. For example, a cell that shows a difference between two tasks performed within the same environment (as reported by Markus *et al.*, 1995), still only fires within a constrained place field in each task, if it fires at all. To say that a cell is sensitive to a non-spatial aspect such as task means that if the cell has a place field under one condition, it may or may not show a place field under the other, and if two cells both show place fields under both conditions, then the spatial relationships between them may change drastically from one condition to the other. Essentially, a cell's place field under one non-spatial condition (in fact whether it has a place field at all) is independent of its field under other non-spatial conditions.

One way to explain this is the multi-map hypothesis: *multiple maps in the hippocampus* [O'Keefe and Nadel, 1978], *active subsets* [Muller and Kubie, 1987], *reference frames* [Wan *et al.*, 1994a, Wan *et al.*, 1994b, Touretzky and Redish, 1996, Redish and Touretzky, 1997a, Redish, 1997], or *charts* [McNaughton *et al.*, 1996, Samsonovich and McNaughton, 1997, Samsonovich, 1997]. All of these authors suggest that the hippocampus includes multiple maps, and that each place cell takes part in one or more of those maps.

There are, however, minor differences among these hypotheses. In this paper, we will concentrate on the two hypotheses which are the most detailed computationally: *charts* and *reference frames*. The multi-chart hypothesis suggests that the maps are critically a hippocampal property arising from internal dynamics of the hippocampus. In contrast, the reference frame hypothesis suggests that the maps arise from interactions of the hippocampus with extrinsic navigational structures such as a neural path integrator.[1]

Barnes *et al.* (1997) have shown an experiment in which the removal and return of an animal to an environment is sufficient to produce a map transition. In this paper, we present an explanation for this result, with simulations. Our account says that errors in a *path integrator reset* process which occur on returning to the environment force a change in reference frame, leading to the appearance of a map transition.

2 The reference frame theory

Over the last few years, we have synthesized a theory of rodent navigation, bringing together ideas from the extensive work done on rodent navigation over the last century and showing how the interaction of several subsystems gives rise to a comprehensive, computational theory of navigation [Touretzky and Redish, 1996, Redish and Touretzky, 1997a, Redish and Touretzky, 1997c]. A complete description of the theory and its correspondence to the experimental literature is given in depth in [Redish, 1997]. Here we will only present a short overview of the theory. Then we will

[1] Path integration is a process that tracks an animal's position as it moves, allowing it to later return to the starting point using only idiothetic cues [Barlow, 1964, Mittelstaedt and Mittelstaedt, 1980, Gallistel, 1990].

compare its explanation for the Barnes *et al.* (1997) experiment with that of the *multi-chart* model [McNaughton *et al.*, 1996, Samsonovich and McNaughton, 1997, Samsonovich, 1997].

Extending O'Keefe and Nadel (1978), the reference frame theory describes navigation as a consequence of four different functional systems:

- **taxon** navigation (direct approach/avoidance of a landmark),
- **praxic** navigation (a sequence of motor actions, driven from an internal sequencing mechanism),
- **locale** navigation (map-based navigation), and
- **route** navigation (chained stimulus-response mechanisms).

It also describes locale navigation in detail as a consequence of an interaction among five spatial representations:

- **local view** (spatial aspects of external landmarks),
- **head direction** (the orientation of the animal in space),
- **path integrator** (the vector home, represented on a canonical map),
- **place code** (a representation of the animal's location in the current reference frame), and
- **goal memory** (allowing trajectory planning).

The anatomical instantiation suggested for these systems is shown in Figure 1.

This theory can explain results from a wide range of methodological paradigms, including single- and multi-cell recording, behavioral manipulations, neuropharmacological manipulations, and lesion studies; it is consistent with anatomical data (see Redish, 1997). We have simulated most aspects of this theory, demonstrating how it can replicate a variety of results, including:

- tracking of head direction by cells in postsubiculum and the anterior thalamic nuclei [Taube *et al.*, 1990, Blair *et al.*, 1997, Taube *et al.*, 1996, see Redish *et al.*, 1996, Goodridge *et al.*, 1997],
- open-field navigation tasks [Collett *et al.*, 1986, Saksida *et al.*, 1995, see Touretzky and Redish, 1996],
- memory consolidation in the Morris water maze [Morris, 1981, Sutherland and Hoesing, 1993, see Redish and Touretzky, 1997c],
- changes in place fields as a consequence of interactions between a consistent entry point and cue card manipulations [Sharp *et al.*, 1990, see Redish and Touretzky, 1996, Redish, 1997], and
- consequences of disorientation in rectangular arenas [Cheng, 1986, Margules and Gallistel, 1988, Gallistel, 1990, see Wan *et al.*, 1994a, Touretzky and Redish, 1996].

All of the simulations are also detailed in [Redish, 1997].

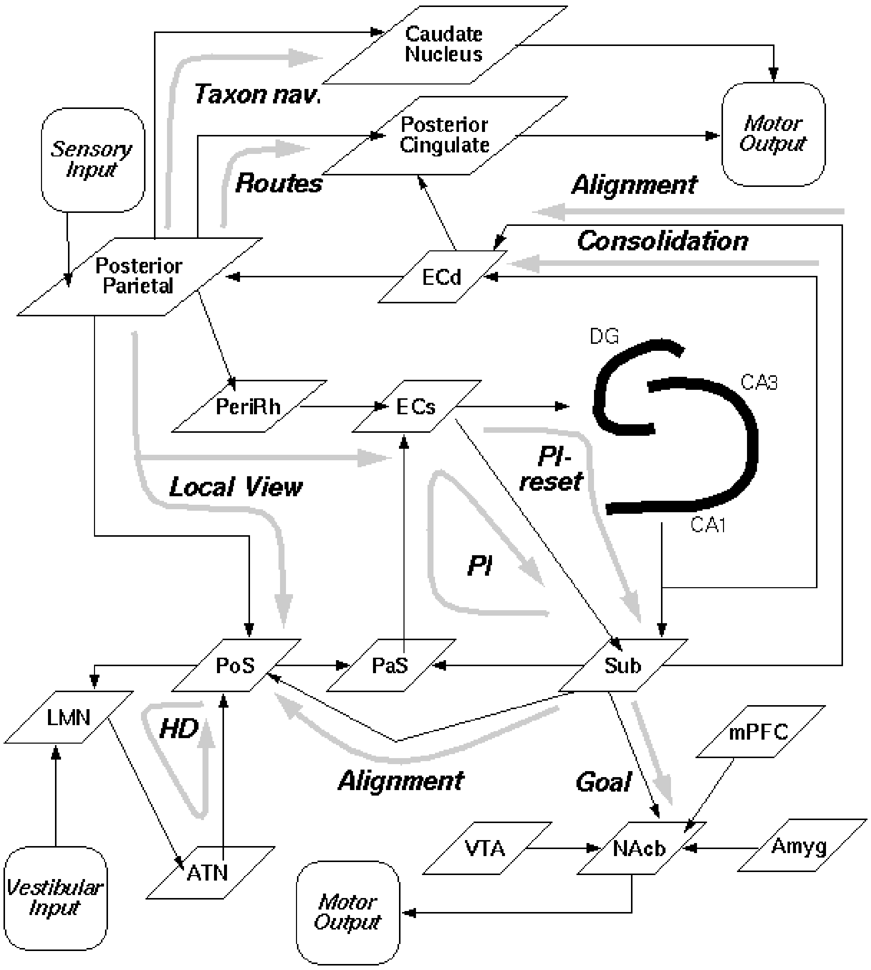

Figure 1: Anatomical realization of a comprehensive model of rodent navigation. From Redish (1997). Amyg: amygdala; ATN: anterior thalamic nuclei; DG: dentate gyrus; CA3,CA1: hippocampus proper; ECs: superficial entorhinal cortex; ECd: deep entorhinal cortex; LMN: lateral mammillary nuclei; mPFC: medial prefrontal cortex. NAcb: nucleus accumbens; PaS: parasubiculum; PeriRh: Perirhinal cortex; PoS: postsubiculum; Sub: subiculum; VTA: ventral tegmental area; HD: Head direction subsystem pathways; PI: Path integration subsystem pathways. Not all anatomical structures or connections are shown. Functional pathways are meant to be indicative only; structures not directly on a labeled pathway may also be involved in that subsystem.

3 The experiment of Barnes *et al.* (1997).

Barnes *et al.* (1997) allowed an animal to walk around a figure-eight maze for 25 minutes. They then removed the animal for one hour, after which the animal was returned to the maze and allowed to walk around for another 25 minutes. During each 25 minute experience, Barnes *et al.* recorded about three dozen place cells simultaneously.

When young animals returned to the environment, they used the same set of place cells to encode location. But when old animals returned to the environment they sometimes used a completely different set of cells. The ensemble correlation between place fields in the two experiences was always high for young animals (approximately 0.7, indicative of a similar representation between experiences), but was bimodal in old animals (sometimes near 0, indicative of a complete remapping, other times near 0.7, indicative of a similar representation). Within a single experience in the environment, place fields were very stable; correlations between the first and second halves of a single run were always high for both old and young animals.

4 Two competing explanations

4.1 The explanation provided by the multi-chart model

Barnes *et al.* (1997, see also McNaughton *et al.*, 1996, Samsonovich and McNaughton, 1997, Samsonovich, 1997) explain the bimodal distribution in old animals as a problem in *selecting the correct cognitive map.*

Their theory includes a set of pre-wired *charts* in the hippocampus, such that the synaptic weight between two cells in hippocampus is inversely proportional to the minimum of their distances across all of the charts. When combined with global inhibition, this produces a local-excitation-global-inhibition network structure. This type of network has a coherent representation of a single location on a single chart as a stable state [Samsonovich and McNaughton, 1997]; any other representation, such as noise biased by extra activity at candidate locations suggested by sensory cues, will be unstable and will settle into a stable state.

The major drawback of this theory is that it requires complex pre-wired connections within the hippocampus. Each place cell needs to be more strongly connected to cells with place fields nearby (in some chart) than to cells with place fields that are distant (in all charts). There is evidence that this connection structure exists after exploration [Wilson and McNaughton, 1994], but the theory requires that the connection structure be in place *before exploration.*

According to this theory, on entering a novel environment, one location on one *chart* (map) will win the competition among competing representations and become the preferred representation for the entry point. As young animals explore the environment, representations of the local view become bound to places on the currently active chart. Then, on a return visit to the environment entering at the same point as before, the local view representation biases the dynamics in the hippocampus so that the same representation of location on the same chart is reinstantiated. In the case of old animals, deficiencies in LTP (see Barnes, 1996, for a review) prevent the local view from becoming as tightly bound to the currently active chart during initial exploration. Thus, according to the multi-chart model, on returning to the environment, old animals experience a much weaker bias to select the same location on the same chart as before.

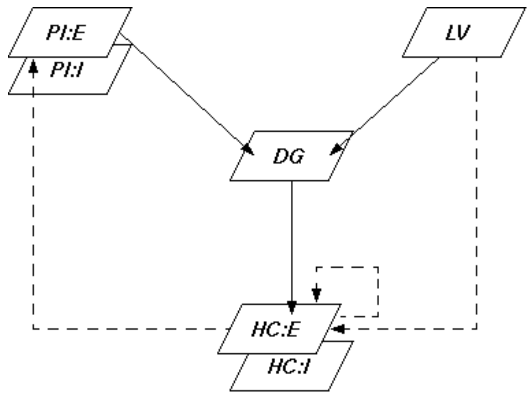

Figure 2: Hippocampal model used to simulate the Barnes *et al.* experiment. During storage, solid lines drive place cell activity and dashed lines show correlational learning; during recall, dashed lines show synaptic transmission and drive place cell activity and path integrator reset. See text for details.

4.2 The explanation provided by the reference frame model

We propose that the phenomenon seen in older animals is not a consequence of pre-wired chart selection within the CA3 population, but rather an interaction between a non-linearity of the path integrator and the orthogonalization properties of dentate gyrus.

The important points for this experiment, drawn from the theory described in Section 2, are:

1. The path integrator is extrinsic to the hippocampus [O'Keefe, 1976, Wan *et al.*, 1994b, Touretzky and Redish, 1996, Redish and Touretzky, 1997a].

2. During normal navigation, place cells require both local view and path integrator input [O'Keefe, 1976, Wan *et al.*, 1994b, Touretzky and Redish, 1996, Redish and Touretzky, 1997a].

3. The dentate gyrus orthogonalizes the combined local view and path integrator inputs [Marr, 1969, McNaughton and Morris, 1987, Rolls, 1989, O'Reilly and McClelland, 1994, Rolls, 1996] so that if either one changes, a new set of place cells is selected.

4. *Path-integrator reset* occurs on re-entry into an environment [Touretzky and Redish, 1996, Redish and Touretzky, 1997a, Redish, 1997, see also Rawlins, 1985, Rotenberg *et al.*, 1996, for similar hypotheses].

Organization of the model. The components required for simulating this experiment are shown in Figure 2. The model includes an extrinsic path integrator (PI), an extrinsic local view (LV), strong random connections from each to the dentate gyrus (DG), and strong random connections from the dentate gyrus to hippocampus (HC). We do not differentiate between CA3 and CA1 in this model, and so HC includes both recurrent connections (as in CA3) and outputs to the path integrator (as in CA1). Both the PI and HC models are composed of excitatory (E) and inhibitory (I) pools.

The path integrator in this model is assumed to consist of a two-dimensional representation of location in which cells show place fields but the fields do not change from environment to environment. Cells in entorhinal cortex and subiculum show these environment-independent place fields

[Quirk *et al.*, 1992, Sharp, 1996]. Following these results, we have suggested that the path integrator consists of a loop between three extra-hippocampal structures: subiculum, parasubiculum, and superficial entorhinal cortex [Redish and Touretzky, 1997a]. The path integrator representation can be updated by offset connections [Zhang, 1996, Samsonovich and McNaughton, 1997].

In addition, we assume that the path integrator has a local-excitation-global-inhibition network structure. This means that the path integrator reset process can occur by assuming the path integrator is initialized with noise and then biased by input from the place cells (which are in turn biased by the local view). This attractor network structure has been extensively studied both in one dimension [Wilson and Cowan, 1973, Amari, 1977, Ermentrout and Cowan, 1979, Kishimoto and Amari, 1979, Kohonen, 1982, Kohonen, 1984, Skaggs *et al.*, 1995, Redish *et al.*, 1996, Zhang, 1996, Redish, 1997] and two [Kohonen, 1982, Kohonen, 1984, Droulez and Berthoz, 1991, Munoz *et al.*, 1991, Arai *et al.*, 1994, McNaughton *et al.*, 1996, Zhang, 1996, Redish and Touretzky, 1997c, Redish, 1997, Samsonovich and McNaughton, 1997, Samsonovich, 1997].

The reference frame model requires that cells in the path integrator be most strongly connected to other cells with nearby place fields, which is similar to the pre-wired connections required in hippocampus in the multi-chart model. However, the reference frame model only requires this connection structure to pre-exist for a single map, located outside the hippocampus, which simplifies the model immensely. The model also includes local excitation (within reference frame) connections within the hippocampus, as does the multi-chart model. However, in the reference frame model this complex connection structure is only assumed to exist after exploration. As has been shown by Muller *et al.* (1991b, 1996, see also Redish and Touretzky, 1997c, Redish, 1997), this connection structure can be learned by random exploration combined with correlational LTP (i.e. Hebbian learning).

Entering a novel environment. When an animal is placed in an environment, we assume that it does not have preconceived path integrator coordinates. The path integrator representation in the model is assumed to initially be random noise. Because the animal has not explored the environment yet, the learnable connections are assumed to have small, uncorrelated random strengths. (The learnable connections, shown by dashed lines in Figure 2, are: LV $\rightarrow$ HC, HC $\rightarrow$ PI, and recurrent connections in HC.) Because these connections are very weak, they do not provide any bias to the settling of the path integrator. Therefore, the path integrator settles to a representation of random coordinates (a "hill" of activation somewhere on the neural sheet) that will serve as the origin or "reference point" for the new reference frame. We call this settling process the *self-localization* or *PI-reset* process.

In contrast to the learnable connnections, the pre-wired connections are assumed to be sparse, and have strong synaptic weights. These connections are indicated by solid lines in Figure 2: LV $\rightarrow$ DG, PI $\rightarrow$ DG, and DG $\rightarrow$ HC.

Activity in the dentate gyrus is a consequence of both the LV and PI representations. In order for a DG cell to fire, it must receive input from both LV *and* PI. Early in the self-localization process, the PI representation is incoherent (i.e. the component neurons show small, random firing rates). This means that early in the self-localization process, DG is effectively silent due to the lack of a coherent representation in PI. This allows the LV $\rightarrow$ HC $\rightarrow$ PI pathway to drive the self-localization process.

In contrast, during navigation, the sparse, strong connections passing through DG drive activity in the hippocampus. Because both LV and PI firing fields are spatially localized, a DG cell will show a high firing rate only in a small, compact portion of the environment, the place field of the cell. Because most of the possible LV $\times$ PI combinations do not occur in an environment, most

DG cells are silent.

Each HC cell receives input from 10–20 DG cells. Activity in one DG cell is sufficient to make the HC cell show a high firing rate. HC cells can therefore have varying numbers of place fields, depending on their specific inputs from DG. In practice, we have found that most HC cells simulated with this model show at most one place field within a reference frame, but occasionally, some cells do show two subfields. Cells with multiple subfields have been reported in real animals (e.g., O'Keefe and Nadel, 1978, Muller *et al.*, 1991a, Wilson and McNaughton, 1994, Markus *et al.*, 1995).

As the animal explores the environment, LTP occurs along the learnable (dashed) connections. We assume that this LTP is Hebbian and rectified at 0, so that synaptic strength can only increase. We do not model LTD.

Returning to a familiar environment. According to the reference frame theory, when young animals return to the environment, LTP has created associations between the LV and HC modules, and between HC and PI. There is a self-localization or PI-reset process each time the animal enters the environment. On a return visit, the learned (dashed) connections that were strengthened by LTP will provide biases to the HC and PI networks. Therefore, upon reentering the environment, the local view representation instantiates a previous representation in hippocampus, and this in turn (via the HC $\rightarrow$ PI) pathway forces the path integrator to reset to the same representation of location as in the young animal's previous experience.

In old animals, however, LTP is deficient (as reviewed by Barnes, 1996) and thus there is little or no bias along the LV $\rightarrow$ HC pathway to reset the hippocampus, and hence via the HC $\rightarrow$ PI pathway, the path integrator, to the same location. Because each DG cell performs a logical *and* function of its path integrator and local view inputs, if the path integrator representation is not reset correctly, this produces a dramatic change in DG representation which will be seen in hippocampus as a low overlap of place codes across different visits to the environment.

Effect of nonlinearity. The attractor network structure hypothesized to underlie the path integrator has an important nonlinearity dependent on where excitatory bias is input into the network. There are four important cases, depending on the location and magnitude of the extra-population input [Redish and Touretzky, 1997b, Redish, 1997, see also Skaggs *et al.*, 1995, Elga *et al.*, 1996, Redish *et al.*, 1996, Zhang, 1996, Samsonovich and McNaughton, 1997, Samsonovich, 1997 for discussions of specific cases]:

1. If an attractor network is in a stable state and receives input (synapsing on excitatory cells) that is peaked at the same position as is currently being represented, then nothing will change. The attractor network will still be in a stable state representing the same position. The overall activity in the attractor network may increase slightly, but the represented position will not change.

2. If the input is offset slightly, then the attractor network will precess until the new representation is centered at the input position.

3. If the input is offset by a large amount but is small in magnitude, it will not be strong enough to affect the current representation, and so the representation will not change.

4. If strong enough input is offset by a large amount, the hill of activiation will jump, i.e. the representation of the current position will disappear and activity will reappear at the offset location.

The effect of this nonlinearity is that if the bias supplied by the HC $\rightarrow$ PI connections is near the position that the path integrator is settling to, it will draw in the representation to match it, whether in strong LTP ("young") or weak LTP ("old") animals. However, if the bias is distant from the position to which the path integrator is settling, the representation will jump only if it is sufficiently strong, i.e. in young animals but not in old animals.

4.3 Similarities and differences

The interpretations of this experiment offered by the multi-chart and reference frame models have some similarities and some crucial differences.

Because place cells are active on initial entry into the environment [Hill, 1978, Austin *et al.*, 1990, Wilson and McNaughton, 1994, Tanila *et al.*, 1997], there must be some pre-wired connections producing place field activity. In the multi-chart model, the charts are pre-wired in CA3. In the reference frame model, pre-wired connections labeled PI $\rightarrow$ DG, LV $\rightarrow$ DG, and DG $\rightarrow$ HC produce place cells with stable fields on initial entry into the environment. The difference between the pre-wired connections hypothesized by the multi-chart model and those hypothesized by the reference frame model is that the latter are initially random.

Because the place cell instability observed by Barnes *et al.* (1997) is bimodal in older animals, there must be some sort of nonlinear process occurring during reentry. In the chart-model, this nonlinearity exists in the competitive dynamics between charts in CA3. In the reference frame model, it is found in the nonlinear settling behavior of the path integrator.

Because the place cell instability observed by Barnes *et al.* (1997) only occurs on entry into the environment, there must be something special about entry into the environment. We explain this by hypothesizing that the path integrator is only reset on entry into the environment. During normal navigation, the path integrator is not reset; it continues to be driven by internal dynamics more than external. But on returning to an environment, the path integrator is reset and external dynamics can have a strong influence.

5 Simulations

Simulations were based on the model shown in Figure 2. All entries to the environment were made at a single location. Each triggered a self-localization process, after which the model was allowed to learn at that location for a short time. The model was then removed from the environment and returned to the same location again, triggering another self-localization.. Each self-localization process began with the path integrator reset to random noise. This cycle of removal and return was repeated 10 times for the simulated young animals with normal LTP, and for simulated old animals with weak LTP.

Because the LV $\rightarrow$ HC and HC $\rightarrow$ PI connections learned the correct mapping in the simulated young animals, the model always reset the path integrator to the same coordinates (Table 1, left). This produced the same PI $\times$ LV association, and thus the same DG and CA3 cells became active. All correlations seen were high, as shown in the left panel of Figure 3.

However, in the simulated "old" animals, the path integrator settled to different coordinates on different tirals (Table 1, right), although the local view had not changed. This changed the PI $\times$ LV association, engendering a change in the DG and thus the CA3 representations. Typical correlations were around 0. But occasionally the path integrator would reset to a location close to its previous value, producing a rare high correlation (> 0.6). See the right panel of Figure 3.

As can be seen in the right half of Table 1, in the last four trials, the simulated old animals always returned to the same reference point. This occured because LTP built up over time. After enough experience the old animals were able to recognize the environment and reset their path integrator representation correctly. The specific timing of the onset of correct reset is strongly dependent on the parameters chosen for the LTP. In the simulation, old animals had LTP that was 5% that of young animals.

In comparing these results to those reported by Barnes *et al.*, we should note that Barnes *et al.* only examined pairs of experiences and so were unable to examine whether old animals eventually settled on a single location. However, Barnes *et al.* tested all of their animals (young and old) in the hidden-platform water maze [Morris, 1981]. They report that early in training, histograms of length of path taken by all animals are bimodal: sometimes they find the platform with a short path and sometimes they require a long path. Barnes *et al.* explain the bimodality as a consequence of being on the correct map or not. If the hippocampus represents the correct location on the correct map, the animal should be able to navigate to the goal; but if the hippocampus has an unfamiliar representation, finding the platform will be much more difficult. As young animals gain experience, they transition to the short-path peak in the bimodal histogram. Old animals also show more short-path trials after four days of training than they do after two, but they continue to be more likely to take a long path than young after four days. In agreement with Barnes *et al.*, we take this as an indication that the old animals are learning to return to the correct map, but they learn much more slowly than young animals.

6 Discussion

These simulations show that the dramatic nonlinearity seen by Barnes *et al.* (1997) does not necessarily imply that there must be pre-wired charts in the hippocampus. A model that includes random pre-wired connections from an external path integrator and an external local view into an orthogonalizing structure (such as dentate gyrus) as well as a single-map pre-wired path integrator is sufficient to produce the observed bimodality in animals with deficient LTP.

This demonstrates why it is important to consider *reference frame* as a property of the entire navigation system and not just the hippocampus. What we are suggesting here is that the old animals are on a different map (different chart, different reference frame) because they are using a different reference point for the path integrator. Because the path integrator representation is changed between experiences, the PI $\times$ LV association changes, and this is why the animal constructs a different reference frame.

The work presented in this chapter does not prove that the pre-wired CA3 hypothesis is wrong, only that it is not a necessary conclusion of the bimodality result reported by Barnes *et al.* (1997). The two theories (that DG is orthogonalizing extrinsic path integrator and local view inputs, and that CA3 is pre-wired to form multiple charts) are not incompatible. Further work is being done to examine the interaction between the two hypotheses.

Acknowledgments

This work was supported by National Science Foundation grant IBN-9631336. The results have also been reported in part in [Redish, 1997]. We thank Carol Barnes, Bruce McNaughton and Jay McClelland for helpful comments and discussions.

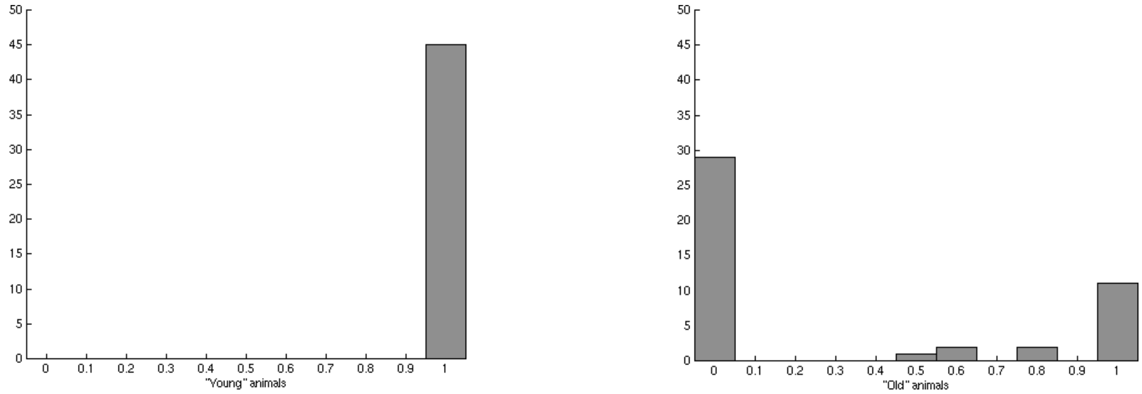

Figure 3: Histograms of correlations between representations of a specific location in the environment after multiple entries into that environment. Simulations were allowed to enter an environment 10 times and the hippocampal representation of a location in the environment was measured. Cross-correlations were made between each pair of experiences in the environment (45 pairs). Plotted is a histogram of the correlations found. (left) Simulations with strong LTP (i.e. "young" animals) always return to the same PI representation, so the correlation is always high. (right) Simulations with weak LTP (i.e. "old" animals) usually return to different PI representations and so usually have very low correlations, but occasionally return to similar PI representations and thus have rare highly correlated pairs.

Entry	"Young" animals	"Old" animals
1	$(120°, 12°)$	$(120°, 12°)$
2	$(120°, 12°)$	$(35°, 131°)$
3	$(120°, 12°)$	$(36°, 113°)$
4	$(120°, 12°)$	$(91°, 63°)$
5	$(120°, 12°)$	$(62°, 109°)$
6	$(120°, 12°)$	$(178°, 326°)$
7	$(120°, 12°)$	$(128°, 11°)$
8	$(120°, 12°)$	$(126°, 12°)$
9	$(120°, 12°)$	$(126°, 12°)$
10	$(120°, 12°)$	$(125°, 12°)$

Table 1: Coordinates represented in the path integrator after each of 10 entries into an environment. Since the path integrator was simulated as a torus, the coordinates are expressed as pairs of angles. All entries were to the same location in the same environment. These coordinates can be taken as indicative of the reference point that would be used for any subsequent navigation.

References

[Amari, 1977] S.-I. Amari. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27:77–87, 1977.

[Arai *et al.*, 1994] K. Arai, E. L. Keller, and J. A. Edelman. Two-dimensional neural network model of the primate saccadic system. *Neural Networks*, 7(6/7):1115–1135, 1994.

[Austin *et al.*, 1990] K. B. Austin, W. F. Fortin, and M. L. Shapiro. Place fields are altered by NMDA antagonist MK-801 during spatial learning. *Society for Neuroscience Abstracts*, 16:263, 1990.

[Barlow, 1964] J. S. Barlow. Inertial navigation as a basis for animal navigation. *Journal of Theoretical Biology*, 6:76–117, 1964.

[Barnes *et al.*, 1997] C. A. Barnes, M. S. Suster, J. Shen, and B. L. McNaughton. Multistability of cognitive maps in the hippocampus of old rats. *Nature*, 388(6639):272–275, 1997.

[Barnes, 1996] C. A. Barnes. Memory changes during normal aging: Neurobiological correlates. In J. L. Martinez, Jr. and R. P. Kesner, editors, *Learning and Memory: A biological view*. Academic Press, New York, 1996. In Press.

[Blair *et al.*, 1997] H. T. Blair, B. W. Lipscomb, and P. E. Sharp. Anticipatory time intervals of head-direction cells in the anterior thalamus of the rat, implications for path integration in the head-direction circuit. *Journal of Neurophysiology*, 1997. In press.

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

[Cohen and Eichenbaum, 1993] N. J. Cohen and H. Eichenbaum. *Memory, Amnesia, and the Hippocampal System*. MIT Press, Cambridge, MA, 1993.

[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.

[Droulez and Berthoz, 1991] J. Droulez and A. Berthoz. A neural network model of sensoritopic maps with predictive short-term memory properties. *Proceedings of the National Academy of Sciences, USA*, 88:9653–9657, 1991.

[Eichenbaum and Cohen, 1988] H. Eichenbaum and N. J. Cohen. Representation in the hippocampus: What do hippocampal neurons code? *Trends in Neurosciences*, 11(6):244–248, 1988.

[Eichenbaum *et al.*, 1987] H. Eichenbaum, M. Kuperstein, A. Fagan, and J. Nagode. Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. *Journal of Neuroscience*, 7(3):716–732, 1987.

[Eichenbaum *et al.*, 1992] H. Eichenbaum, T. Otto, and N. J. Cohen. The hippocampus – what does it do? *Behavioral and Neural Biology*, 57:2–36, 1992.

[Eichenbaum, 1996] H. Eichenbaum. Is the rodent hippocampus just for "place?". *Current Opinion in Neurobiology*, 6:187–195, 1996.

[Elga *et al.*, 1996] A. N. Elga, A. D. Redish, and D. S. Touretzky. A model of the rodent head direction system. *CNS*96 (Fifth Annual Computational Neuroscience Meeting).*, 1996.

[Ermentrout and Cowan, 1979] B. Ermentrout and J. Cowan. A mathematical theory of visual hallucination patterns. *Biological Cybernetics*, 34:137–150, 1979.

[Gallistel, 1990] C. R. Gallistel. *The Organization of Learning.* MIT Press, Cambridge, MA, 1990.

[Goodridge *et al.*, 1997] J. P. Goodridge, A. D. Redish, D. S. Touretzky, H. T. Blair, and P. E. Sharp. Lateral mamillary input explains distortions in tuning curve shapes of anterior thalamic head direction cells. *Society for Neuroscience Abstracts*, 23, 1997. In press.

[Gothard *et al.*, 1996] K. M. Gothard, W. E. Skaggs, K. M. Moore, and B. L. McNaughton. Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *Journal of Neuroscience*, 16(2):823–835, 1996.

[Hampson *et al.*, 1993] R. E. Hampson, C. J. Heyser, and S. A. Deadwyler. Hippocampal cell firing correlates of delayed-match-to-sample performance in the rat. *Behavioral Neuroscience*, 107(5):715–739, 1993.

[Hill, 1978] A. J. Hill. First occurrence of hippocampal spatial firing in a new environment. *Experimental Neurology*, 62:282–297, 1978.

[Kishimoto and Amari, 1979] K. Kishimoto and S. Amari. Existence and stability of local excitations in homogenous neural fields. *Journal of Mathematical Biology*, 7(4):303–318, 1979.

[Kohonen, 1982] T. Kohonen. Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69, 1982.

[Kohonen, 1984] T. Kohonen. *Self-Organization and Associative Memory.* Springer-Verlag, New York, 1984.

[Kubie and Ranck, 1983] J. L. Kubie and J. B. Ranck, Jr. Sensory-behavioral correlates in individual hippocampus neurons in three situations: Space and context. In W. Seifert, editor, *Neurobiology of the Hippocampus*, chapter 22, pages 433–447. Academic Press Inc., New York, 1983.

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

[Markus *et al.*, 1995] E. J. Markus, Y. Qin, B. Leonard, W. E. Skaggs, B. L. McNaughton, and C. A. Barnes. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *Journal of Neuroscience*, 15:7079–7094, 1995.

[Marr, 1969] D. Marr. A theory of cerebellar cortex. *Journal of Physiology*, 202:437–470, 1969. Reprinted in Marr (1991).

[Marr, 1991] D. Marr. *From the Retina to the Neocortex: Selected Papers of David Marr.* Birkhäuser, Boston, 1991. Edited by L. M. Vaina.

[McNaughton and Morris, 1987] B. L. McNaughton and R. G. M. Morris. Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences*, 10(10):408–415, 1987.

[McNaughton *et al.*, 1996] B. L. McNaughton, C. A. Barnes, J. L. Gerrard, K. Gothard, M. W. Jung, J. J. Knierim, H. Kudrimoti, Y. Qin, W. E. Skaggs, M. Suster, and K. L. Weaver. Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *Journal of Experimental Biology*, 199(1):173–186, 1996.

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

[Morris, 1981] R. G. M. Morris. Spatial localization does not require the presence of local cues. *Learning and Motivation*, 12:239–260, 1981.

[Muller and 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.*, 1991a] 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.

[Muller *et al.*, 1991b] R. U. Muller, J. L. Kubie, and R. Saypoff. The hippocampus as a cognitive graph. *Hippocampus*, 1(3):243–246, 1991.

[Muller *et al.*, 1996] R. U. Muller, M. Stead, and J. Pach. The hippocampus as a cognitive graph. *Journal of General Physiology*, 107(6):663–694, 1996.

[Munoz *et al.*, 1991] D. P. Munoz, D. Pélisson, and D. Guitton. Movement of neural activity on the superior colliculus motor map during gaze shifts. *Science*, 251:1358–1360, 1991.

[O'Keefe and Nadel, 1978] J. O'Keefe and L. Nadel. *The Hippocampus as a Cognitive Map*. Clarendon Press, Oxford, 1978.

[O'Keefe, 1976] J. O'Keefe. Place units in the hippocampus of the freely moving rat. *Experimental Neurology*, 51:78–109, 1976.

[O'Reilly and McClelland, 1994] R. C. O'Reilly and J. L. McClelland. Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus*, 4(6):661–682, 1994.

[Quirk *et al.*, 1992] G. J. Quirk, R. U. Muller, J. L. Kubie, and J. B. Ranck, Jr. The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *Journal of Neuroscience*, 12(5):1945–1963, 1992.

[Rawlins, 1985] J. N. P. Rawlins. Associations across time: The hippocampus as a temporary memory store. *Behavioral and Brain Sciences*, 8:479–496, 1985. See also commentary and response, pages 497–528.

[Redish and Touretzky, 1996] A. D. Redish and D. S. Touretzky. Modeling interactions of the rat's place and head direction systems. In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems 8*, pages 61–71, 1996.

[Redish and Touretzky, 1997a] A. D. Redish and D. S. Touretzky. Cognitive maps beyond the hippocampus. *Hippocampus*, 7(1):15–35, 1997.

[Redish and Touretzky, 1997b] A. D. Redish and D. S. Touretzky. Implications of attractor networks for cue conflict situations. *Society for Neuroscience Abstracts*, 23:1601, 1997.

[Redish and Touretzky, 1997c] A. D. Redish and D. S. Touretzky. The role of the hippocampus in solving the Morris water maze. *Neural Computation*, 1997. In press.

[Redish *et al.*, 1996] A. D. Redish, A. N. Elga, and D. S. Touretzky. A coupled attractor model of the rodent head direction system. *Network: Computation in Neural Systems*, 7(4):671–685, 1996.

[Redish, 1997] A. D. Redish. *Beyond the Cognitive Map: Contributions to a Computational Neuroscience Theory of Rodent Navigation.* PhD thesis, Carnegie Mellon University, Pittsburgh PA, 1997.

[Rolls, 1989] 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*, chapter 8, pages 125–159. Addison-Wesley, Reading, MA, 1989.

[Rolls, 1996] E. T. Rolls. A theory of hippocampal function in memory. *Hippocampus*, 6:601–620, 1996.

[Rotenberg *et al.*, 1996] A. Rotenberg, M. Mayford, R. D. Hawkins, E. R. Kandel, and R. U. Muller. Mice expressing activated CaMKII lack low frequency LTP and do not form stable place cells in the CA1 region of the hippocampus. *Cell*, 87:1351–1361, 1996.

[Saksida *et al.*, 1995] L. M. Saksida, A. D. Redish, C. R. Milberg, S. J. Gaulin, and D. S. Touretzky. Landmark-based navigation in gerbils supports vector voting. *Society for Neuroscience Abstracts*, 21, 1995.

[Samsonovich and McNaughton, 1997] A. Samsonovich and B. L. McNaughton. Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, 17(15):5900–5920, 1997.

[Samsonovich, 1997] A. V. Samsonovich. *Attractor map theory of the hippocampal representation of space.* PhD thesis, Graduate Interdisciplinary Program in Applied Mathematics, University of Arizona, 1997.

[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, 1990.

[Sharp, 1996] P. E. Sharp. Subicular cells generate similar firing patterns in two geometrically and visually distinctive environments: Comparison with hippocampal place cells. *Behavioral Brain Research*, 1996. In press.

[Skaggs *et al.*, 1995] W. E. Skaggs, J. J. Knierim, H. S. Kudrimoti, and B. L. McNaughton. A model of the neural basis of the rat's sense of direction. In G. Tesauro, D. S. Touretzky, and T. K. Leen, editors, *Advances in Neural Information Processing Systems 7*, pages 173–180. MIT Press, 1995.

[Sutherland and Hoesing, 1993] R. J. Sutherland and J. M. Hoesing. Posterior cingulate cortex and spatial memory: A microlimnology analysis. In B. A. Vogt and M. Gabriel, editors, *Neurobiology of Cingulate Cortex and Limbic Thalamus: A Comprehensive Handbook*, pages 461–477. Birkhauser, Boston, 1993.

[Tanila *et al.*, 1997] H. Tanila, P. Sipilä, M. Shapiro, and H. Eichenbaum. Brain aging: Impaired coding of novel environmental cues. *Journal of Neuroscience*, 17(13):5167–5174, 1997.

[Taube *et al.*, 1990] 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.*, 1996] J. S. Taube, J. P. Goodridge, E. J. Golob, P. A. Dudchenko, and R. W. Stackman. Processing the head direction cell signal: A review and commentary. *Brain Research Bulletin*, 40(5/6):477–486, 1996.

[Thompson and Best, 1989] L. T. Thompson and P. J. Best. Place cells and silent cells in the hippocampus of freely-behaving rats. *Journal of Neuroscience*, 9(7):2382–2390, 1989.

[Touretzky and Redish, 1996] D. S. Touretzky and A. D. Redish. A theory of rodent navigation based on interacting representations of space. *Hippocampus*, 6(3):247–270, 1996.

[Wan *et al.*, 1994a] H. S. Wan, D. S. Touretzky, and A. D. Redish. A rodent navigation model that combines place code, head direction, and path integration information. *Society for Neuroscience Abstracts*, 20:1205, 1994.

[Wan *et al.*, 1994b] 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 Earlbaum Associates, Hillsdale NJ, 1994.

[Wiener, 1993] S. I. Wiener. Spatial and behavioral correlates of striatal neurons in rats performing a self-initiated navigation task. *Journal of Neuroscience*, 13(9):3802–3817, 1993.

[Wilson and Cowan, 1973] H. R. Wilson and J. D. Cowan. A mathematical theory of the functional dynamics of cortical and thalamic tissue. *Kybernetik*, 13:55–80, 1973.

[Wilson and McNaughton, 1994] M. A. Wilson and B. L. McNaughton. Reactivation of hippocampal ensemble memories during sleep. *Science*, 265:676–679, 1994.

[Zhang, 1996] K. Zhang. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience*, 16(6):2112–2126, 1996.