



Think Like a Rat

Charles C. Kemp
November 15, 2001

Abstract

Since 1971 when the results of single electrode recordings of neurons with “place fields” were first documented, scientists have produced an enormous body of research in an attempt to understand these small but significant components of the rat brain. This research has confirmed that after a rat has gained familiarity with a simple test environment, a neuron will respond strongly and selectively to the rat’s presence within a small part of that environment. But despite 30 years of research, remarkably little is known about the computations that produce and use the output of these “place cells”.

First, a short review of scientific results related to place cells and the rat hippocampus highlights the rewards that might come from understanding these computations. Then a review and analysis of two recent computational models of rat navigation that strongly involve place cells provides a view of the current state of the art in this research. Among other observations, a part of the model by Foster, Morris and Dayan (2000), designed to solve the problem of learning from a single example, is shown to rely on questionable assumptions that might lead to failure of the model on real data. As an example of the under constrained space of solutions given current experimental data, another possible solution to the problem of single example learning is proposed. The model by Arleo & Gerstner (2000) is shown to conflict with several results from recent experiments on rats. Although both models have flaws, elements from each of them supply ideas worthy of consideration, which are isolated, compared and discussed.

After this analysis of the two models, a quick review of options for future research into computational models of place cells is presented with an eye towards accelerating science’s progress. Within this review several experimental techniques are looked at including recent developments in the use of multi-electrodes pioneered by Wilson and McNaughton (1993).

1. Introduction

With good reason, scientists have studied rats for over a century in an effort to better understand mammalian intelligence. Neuroscience has shown striking similarities between the brains of rats and humans, yet rats are much simpler in many ways that would seem to make them easier to understand. One quantifiable measure of this relative simplicity is that on average a rat’s brain weighs 2g while a human’s weighs 1400g (Chudler, 2001). Moreover, their small size, rapid breeding, and general ease of maintenance, has helped rats become a preferred animal for laboratory experiments, a preference that has lead to vast knowledge of their biology.

Over most of the last hundred years, rat navigation has been a common method of exploring rat intelligence. This is understandable, given that rat locomotion is the most easily quantified and observed output of rat intelligence. At the beginning of the 20th century research in the brain sciences regularly involved running rats through mazes in attempts to investigate properties of intelligence such as learning and memory, but in spite of a large number of experiments progress was very limited. (Redish, 1999).




Figure 1. A modern recording and representation of hippocampal place cells from Wilson & McNaughton (1993). Each box represents the average firing of a single hippocampal neuron as a function of the rat's position within a square environment. Red represents a higher average firing rate and blue represents quiescence. The three mostly red boxes are probably non-pyramidal neurons called interneurons or theta neurons. These functions are typically referred to as a neuron's receptive field. The functions with small, localized regions of sensitivity are typically called "place fields" and their associated neurons are often called "place cells". (Wilson & McNaughton, 1993)

In 1971 exciting new information was discovered about the rat brain that would come to greatly increase the stakes for understanding rat navigation. In this year, O'Keefe & Nadel published results from single electrode recordings in the rat hippocampus that indicated that pyramidal neurons in the region CA1 were sensitive to the rat's current location and could be characterized by surprisingly intuitive "place fields", see figure 1 (O'Keefe, 1999). This result is astounding since the recordings were from single neurons out of approximately 150 million neurons total in the rat brain and half a million pyramidal neurons in the hippocampus (Chudler, 2001). This discovery is even more remarkable considering the placement of the hippocampus within the architecture of the rat brain. Anatomical and experimental evidence indicate that the hippocampus sits in the middle of computations performed by the rat brain, receiving input from cortical areas that are far removed from initial sensory processing, and outputting information to several areas that precede motor control considerably.

The discovery of these place cells within the hippocampus very much enhances their significance. The hippocampus is one of the most well studied brain structures of the mammalian brain with closely related analogues in species including rats and humans. Anatomically it was isolated in the 1800's because of its very distinctive organization and regular structure (Redish, book). In addition, it is a brain structure of great importance within humans, where among other roles it is crucial to spatial and episodic memory. A famous example of this is the patient H.M. whose hippocampi were bilaterally removed and has been unable to explicitly recall any of the experiences he has had since the operation.

Besides anatomical evidence for similar functions performed by the rat and human hippocampus, many experiments have convincingly demonstrated the use of the human hippocampus in navigation and spatial reasoning (Kandel et al, 2000). People with damage to their left hippocampus show clear deficits in navigation tasks and other spatially oriented activities (Spiers et al, Brain 2001). Recent studies with people using virtual reality environments have related the left human hippocampus to both navigation and spatial reasoning through imaging techniques and tests on people with hippocampal damage (Maquire et al, 1998)(Burgess et al, 2001). Single electrode recordings of hippocampal cells within monkeys have demonstrated the existence of view sensitive cells that may be closely related to place cells in rats (Araujo et al, 2001).



Figure 2. The yellow region is a 3d reconstruction of the hippocampus volume within a rat brain. The picture of the rat brain on the right is taken from above the rat's head looking down, where the front of the rat's head would be facing towards the top of the page. The gray surface that sits above the hippocampus in the right image is the cortex (Fiala et al 2001)

Some experiments also indicate that the rat hippocampus may represent more general contextual information akin to human episodic memory, along with the rats' current location (Wood et al, 1999). Among many other fascinating discoveries made about the hippocampus is its activity during sleep. In rats, multi-electrode recordings of the hippocampal place cells have shown that activity patterns due to the locations the rat visited while awake are to some extent reflected in place cell activity during slow-wave sleep (Wilson & McNaughton, 1994). More recent research has even demonstrated that some sequences of patterns may be replayed during REM sleep, which could be indicative of the consolidation of memories into the neocortex (Louie & Wilson, 2001).




Figure 3. Hippocampal cross section of a neonatal kitten originally published in 1903, with silver chromate staining that preferentially acts on the pyramidal neurons within CA1. Notice the conservation of hippocampal structure between a rat and a cat (Kandel & Schwartz, 2000).

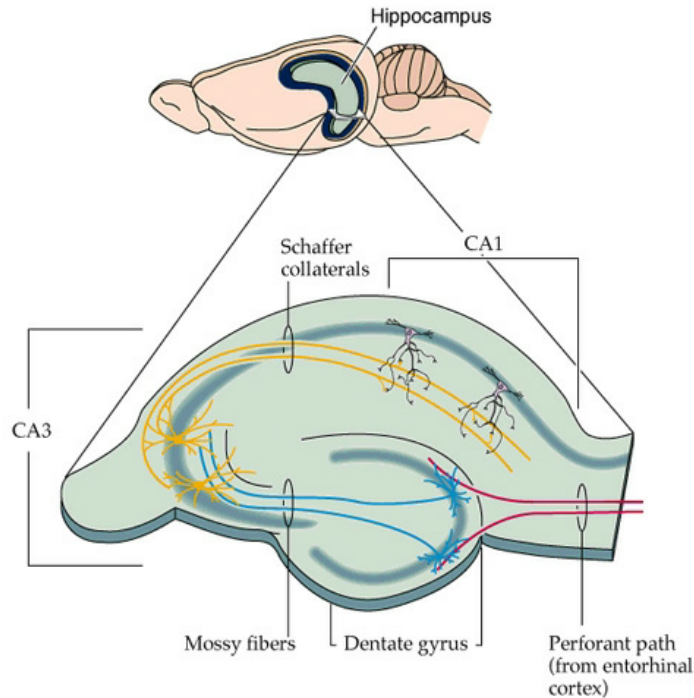


Figure 4. This shows a canonical hippocampal slice. The perforant path projects to the dentate gyrus (DG) from the entorhinal cortex (EC). The DG projects via mossy fibers to CA3 and CA3 projects by the schaffer collaterals to CA1. Subsequently CA1 projects back to the EC as well as to various subcortical structures. (Bulman-Fleming,2000)

As one would expect this history of exciting studies of the hippocampus has led to many theories about the computations performed by the hippocampus. Unfortunately, many of these theories are not explicit about the computations performed. Likewise, even when they are computationally explicit, testable predictions are often difficult to construct. For the rat hippocampus, the two most prominent overarching theories are “The Cognitive Map” hypothesis of O’Keefe (1999) and the hypothesis that the hippocampus serves as a type of general memory system for which Eichenbaum (1996) is a major proponent.

Place cells within the rat hippocampus can serve as a well-grounded point from which to attack these broader issues of hippocampal function. Since, place cells are activated by rat locomotion to different locations within an environment, navigation is a reasonable task to study. Moreover, as Redish points out, scientists, mathematicians, and engineers have already developed a very sophisticated understanding of navigation because of its crucial role in human activities. As such, navigation can serve as a useful domain for the study of biological intelligence, since many of the theoretical tools necessary for objectively describing the tasks involved already exist (Redish, 1999)(Arleo, 2000).

Currently, the majority of computational models involving hippocampal place cells, are based on a coarse scale of tissue organization, if they are related to the neural hardware at all. For the most part the models that do make analogies to real neural tissue are related to large-scale modular structures in the realm of millimeters and the major interconnections among them. These tissue areas are usually structures that can be readily separated through standard techniques of neuroanatomy, such as staining and dissection, see figure 3. Figure 4 shows the main hippocampal regions referenced by these computational models, as well as some of their interconnections. Typically recordings of place cells are made within CA1, although less reliable place cells have been found in many regions within and surrounding the hippocampus.

1.1 Navigation

Experiments have implicated the use of place cells in some forms of rat navigation. Commonly researchers divide navigation into five distinct types, although several other taxonomies have been proposed. The first type, is just random motion. The second type, called taxon navigation, requires the creature to orient towards a stimulus and approach it. The third type, called praxic navigation, is less clear but predominantly relates to fixed sequences of actions performed using self-motion information and minimal feedback from the environment. The fourth and fifth types of navigation are the least clear. The fourth type, called route navigation, seems to refer to chains of fixed action sequences that are triggered by simple stimuli. While the fifth type, called locale navigation, assumes that the creature has a cognitive map that allows the creature to plan a path between its current location and a goal location when it knows where those locations are on the map. (Redish, 1999) (Arleo, 2000) Essentially these classes of navigation serve as vague computational theories that help researchers construct experiments. It seems unlikely that all of these clear divisions exist within the computational systems of animals. The development of more explicit computational models should help enumerate further possibilities. Given this classification, researchers have found that place cells are at least required for locale navigation.

Both Foster et al (2000) and Arleo & Gerstner (2000) propose computational models of rat navigation that strongly involve place cells and reinforcement learning. The next two sections will review these state of the art computational models, after which they will be compared with one another and analyzed.

2. The Model of Foster, Morris and Dayan

The model of Foster et al proposes navigational computations that would make use of the output from hippocampal place cells. Their model does not explore the computations that would be required to produce the output of the place cells; instead they assume an idealized form for the population of place cells. Their model uses these idealized outputs to perform two distinct types of navigation. The first type of navigation assumes a primarily static environment in which rewards are associated with particular locations. Based on this assumption they learn both the value of being at a particular location and the value of taking various actions from each location. The second type of navigation learns metric coordinates for each location in the environment and then simply subtracts these coordinates from the coordinates of a goal in order to determine where the rat should move.

2.1 The Control Functions

The authors model the receptive field f_i of each place cell i by a 2 dimensional Gaussian centered on a location s_i within the environment.

$$f_i(p) = e^{-\frac{\|p-s_i\|^2}{2\sigma^2}}$$

The standard deviation $\sigma = 0.16m$ is constant for all of the Gaussians and their centers are densely and uniformly distributed over the environment at about 157 place fields per m^2 , which causes substantial overlap in their receptive fields, see figure 5.




Figure 5. This cross-section shows the spacing of the radial basis functions used by Foster et al

The two proposed methods of navigation use these receptive fields as a regularizing basis set with which they create several functions. Each function g outputs a scalar value when given a position as input and is formed by a simple linear combination of the receptive fields, which serve as the basis vectors. We can denote this simply by taking the dot product of a weight vector w with a function $f(p)$, which has a vector output whose elements are the basis functions $f_i(p)$.

$$g(p) = w^T f(p)$$

2.1.1 Static Navigation

The first method of navigation, which for the rest of this paper will be called “static navigation”, uses a total of nine functions one of which, $c(p)$, represents an estimate of the expected value, $v(p)$, of being at the input location while the other eight, $a_j(p)$, $j=1\dots 8$, specify the relative value of moving in one of eight directions from the input direction. At each time step, based on the output of these eight functions, each direction has a probability $P[j]$ of being selected as the next command. The probabilities are designed to force the rat to explore the environment.

$$P[j] = \frac{e^{(2a_j)}}{\sum_i e^{(2a_i)}}$$

2.1.2 Metric Navigation

The second method of navigation they propose, which for this paper will be called “metric navigation”, uses two functions, $x(p)$ and $y(p)$, that together associate an x and y coordinate with each position. At each time step this navigation system subtracts $x(p)$ and $y(p)$ from a goal position and then moves the rat in the direction of the resulting vector, unless it is already at the goal position in which case it erases the goal and searches randomly. These two methods of navigation are then combined in so that the second method of navigation gradually takes control of the rat away from the first method of navigation.

2.2 Learning the Functions

Given this framework for control, the challenge for the system is to determine good weight vectors for the 11 functions described above, so that the place field output can be converted into good decision functions.

Their model uses standard algorithms for reinforcement learning to iteratively find weight vectors while the simulated rat is exploring the environment.

2.2.1 Static Navigation Learning

The first navigation method uses the actor critic method of temporal difference (TD) learning, with $c(p)$ as the critic and $a_j(p)$ as the actor. At each time step, $c(p)$ serves as the current estimate of the ideal value function $v(p)$, where p_t is the position at time t , R_t is the reward received at time t , and $\langle \cdot \rangle$ signifies the mean over all trials.

$$v(p_t) = \langle R_t + \gamma R_{t+1} + \gamma^2 R_{t+2} + \dots \rangle$$

Due to their assumption of exponentially discounted rewards a recursion relationship holds for $v(p_t)$.

$$v(p_t) = \langle R_t \rangle + \gamma v(p_{t+1})$$

One trick from TD learning is to use $c(p)$'s deviation from this recursive constraint by $c(p)$, as an error measure by which to train $c(p)$.

$$\delta_t + c(p_t) = R_t + \gamma c(p_{t+1})$$

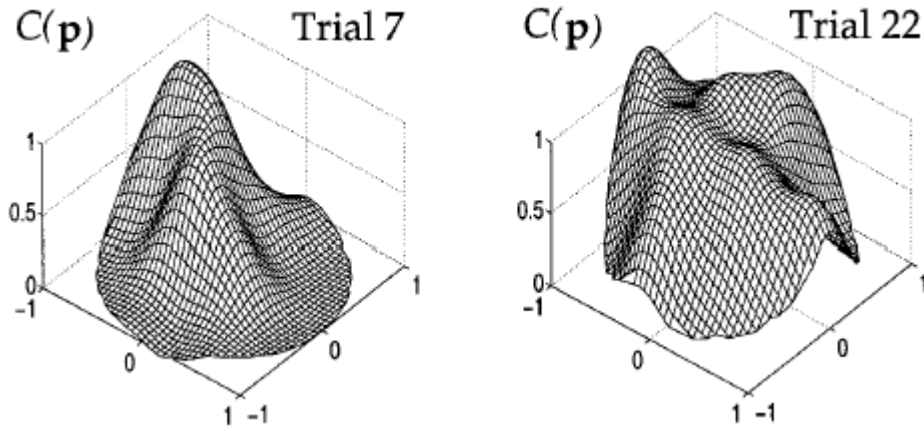


Figure 6. Shows an example of $c(p)$ learning to represent the ideal value function for a reward placed in the upper corner of the environment. (Foster et al, 2000)

Gradient descent on $c(p_t)$ using the error δ_t and a learning rate of α to adjust the associated weight vector w_c so as to lower the error and make $c(p_t)$ more closely approximate $v(p_t)$, see figure.

$$\begin{aligned} \nabla[c]_{w_c} &= f \\ \Delta w_c &= \alpha \delta_t f(p_t) \end{aligned}$$

With a fixed set of action functions, typically called the policy, and some additional easily met conditions $c(p)$ would converge to $v(p)$ (Sutton & Barto, 1998). In this model however, the functions that create the actor are being changed at the same time as the value function. Another method from TD learning is used to improve the actions selected by the action functions. Specifically, δ_t serves as a reinforcement signal to the action functions that informs them about the quality of their most recent action. If $\delta_t > 0$ the actual reward

received from executing the previous action was larger than expected by the critic, and consequently this choice of action from the previous state should be encouraged. Similarly, if $\delta_t < 0$ the actual reward received from executing the previous action was smaller than expected by the critic, and consequently this choice of action from the previous state should be discouraged. Finally, if $\delta_t = 0$ the reward was exactly as expected by the critic so no change should be made to the actor. Due to $P[j]$, increasing the output value of a particular action function increases its probability of selection, while decreasing its output lowers its probability of selection. Consequently, after an action j is taken its weight vector w_j can be beneficially adjusted by changing it along its gradient in proportion to δ_t , while leaving the other seven action functions the same. This is a form of hebbian learning since the action function will only be changed if it was just used. The resulting equations for weight updates are of the exact same form used to adjust the critic's weight vector.

2.2.2 Metric Navigation Learning

The second form of navigation adjusts the weight vectors, w_x and w_y , for its two functions using a different strategy. In this case, the error signal is constructed based on the difference between self-motion information that tells the system how far the rat traveled in the last time step and an estimation of traveled distance by the coordinate functions. This is clever, because the resulting form is identical to TD learning with $\gamma=1$, where Δx_{sm} is the estimated change in position from self-motion and ϵ_t is the error signal analogous to δ_t (Sutton & Barto, 1998).

$$\begin{aligned}\epsilon_t^x + x(p_t) &= \Delta x_{sm} + x(p_{t+1}) \\ \epsilon_t^y + y(p_t) &= \Delta y_{sm} + y(p_{t+1})\end{aligned}$$

They use TD(λ) to train $x(p)$ and $y(p)$, which is a more general form of TD learning that allows them to use longer histories of errors to adapt $x(p)$ and $y(p)$ (Sutton & Barto, 1998). This helps make the coordinate system learning more efficient. Its need is clear when one considers that all positions along a dimension are dependent on one another. This learning is designed to enforce global consistency, which implies the propagation of information across the entire extent of the state space. Otherwise, without sufficient training some domains of disagreement could remain. The resulting training equations for w_x and w_y with a learning rate of κ are as follow:

$$\begin{aligned}\Delta w_x &= \kappa \epsilon_t^x \sum_{k=1}^t \lambda^{t-k} f_i(p_k) \\ \Delta w_y &= \kappa \epsilon_t^y \sum_{k=1}^t \lambda^{t-k} f_i(p_k)\end{aligned}$$

Creating a variable that represents the value of using the coordinate based navigation system instead of the actor combines these two methods of navigation. When a goal is in memory the output of the coordinate based navigation system is acted upon with some probability determined by the value of the variable relative to the eight action functions. When its action is selected δ_t multiplied by a learning rate is added to the current value of the variable. Essentially this variable is a surrogate action function with a weight vector with only one element and an input that is fixed at 1, that serves to represent the option of using the coordinate based navigation system.

3. The Model of Arleo & Gerstner

Arleo proposes a large and complicated model that both navigates based on the output of place cells and computes the output of place cells based on raw sensor data.

3.1 Static Navigation

The navigation side of the model is very similar to the non-coordinate based navigation system in Foster et al's model. Both models use reinforcement learning to incrementally adjust the weight vectors of simple action functions that linearly combine the outputs from the place cells. Likewise, both models take advantage of the broadly tuned and highly redundant place cell representation as a regularizing basis set. In addition, although Arleo & Gerstner tested their place cell model with a robot, they only evaluated the navigation algorithm in simulation. Unfortunately, they do not clearly state the place field representation they used during these simulations

One difference worth noting between the first part of Foster et al's model and Arleo & Gerstner's navigation model, is Arleo & Gerstner's use of Q-learning instead of an actor-critic algorithm. Essentially this standard reinforcement learning algorithm only uses an actor and assumes that at each state the value function, previously represented by the critic, can be approximated by the output of the actor function with the maximum value at that state.

$$c(p) = \text{Max}_j(a_j(p))$$

This leads to the following equation for the error signal that is used to update the action function responsible for the action at time t.

$$\delta_t + a_j(p_t) = R_t + \gamma \text{Max}_j(a_j(p_{t+1}))$$

Another minor difference in Arleo & Gerstner's implementation of reinforcement learning is the use of TD(λ) with its associated memory trace rather than TD(0).

Arleo & Gerstner make two modifications to this simple navigation system. First, when multiple tasks exist, several actors that are each sensitive to a different type of reward are used simultaneously, with one actor controlling the rat's actions in the current task while the other actors opportunistically learn. Second, in order to more efficiently learn when a reward is moved, Arleo & Gerstner added a function that can cause the current actor to reset. Essentially, a hebbian learning algorithm was used to train a simple linear function to expect rewards at locations. When this function indicates a strong failure to attain expected results for several time steps, the current actor is reset, forgetting all that it has learned. Without this workaround, the system would take an extremely long time to learn to navigate to the new reward.

3.2 Place Cells




Figure 7. Arleo's (2000) place fields. The top row gives examples of sEC place fields. The right three show examples of deviant place fields with multiple peaks and less locality, which comprise 10% of the sEC place fields using the gabor filters. The bottom row shows examples of CA1 place fields that combine output from path integration and vision.

A large part of Arleo & Gerstner's model is devoted to the computation of place cell responses using vision and path integration information as input. Arleo & Gerstner's model attempts to use these two sources of location information in a complementary way. After enough time has passed for the path integrator to accumulate significant error, the vision system is used to find a visually distinctive location at which the path integrator can be reset. On the other hand, visually ambiguous locations are distinguished by the output of the path integrator.

The vision system is used to construct a set of place sensitive cells that are solely driven by visual information. Likewise, path integration information drives its own set of place sensitive cells. These two sets of place sensitive cells are then combined to form a model of the hippocampal place cells in the CA1 region. Technically Arleo & Gerstner claim to model CA3-CA1 cells, but this review will stick with the convention of only referring to CA1 cells, since the difference here is unimportant. Arleo & Gerstner hypothesize that these two areas of preliminary place sensitive cells are located in the entorhinal cortex, which as shown in figures 4 & 5, is a cortical region adjacent to the hippocampus that provides much of the input to the hippocampus. They refer to the visually driven cells as sEC cells, for superficial entorhinal cortex.

3.2.1 Visually Driven sEC Place Cells (pre-CA1)




Figure 8. The robot and an example image from the test performed by Arleo (2000)

The Arleo & Gerstner paper describes results from a linear camera carried by a robot in a contrived environment with bar code walls. Within Arleo's thesis, he presents newer results that use a regular gray-scale camera carried by a robot within an apparently unaltered laboratory environment. The following description will only refer to the later work, although the work in the paper has very similar structure.

When the robot enters a location that it believes to be new or inadequately represented, it first collects four images $I_j, j = 1...4$, one in each of the cardinal directions as determined by a system which maintains a globally consistent representation of the robot's orientation. This orientation system is based on a model of head direction cells found in rats (Arleo, 2000). Next, it applies a linear multi-scale filter bank to the images and calculates the magnitudes of the result. A subset of the resulting filter responses, r_j , are then selected for further use, in this case by a log-polar sub-sampling strategy. At this point a snapshot cell is assigned to each of the four resulting filter response vectors, r_j . From then on each snapshot cell tests incoming images to see how similar they are to the image with which the snapshot cell was initialized. Snapshot cells do this by putting the Euclidean distance between their representative response vector r_j and the response vector from the incoming image through a 1D Gaussian. As such, a snap shot cell is a radial basis function centered on the vector, r_j , on which it was initialized.

Finally, these four snap shot cells are assigned to a visually driven sEC place cell, b , which calculates a weighted average of their responses using a four element weight vector, w .

$$b(s) = \frac{w^T s}{\sum_i w_i}$$

After being initialized to random weights between 0 and 1, w is incrementally updated with a simple hebbian learning rule.

$$\Delta w_j = (s_j - w_j)b(s)$$

As shown in the figure 7, these sEC cells apparently do a good job of selectively responding to locations within the environment as long as the global orientation of the robot is maintained.

3.2.2 Path Integration Place Cells (PI cells) (pre-CA1)

Arleo's model for the path integration subsystem explicitly integrates the robot's wheel turns and then represents the resulting position with a set of 2d Gaussian receptive fields. This is the same transformation of position, $f(p)$, that Foster's model uses to simulate the CA1 place fields, although the specific spatial distribution, s_i , and standard deviation, σ , may differ.

3.2.3 CA1 Place Cells

As with the sEC cells, when the robot enters a location that it believes to be new or inadequately represented, it creates a new CA1 place cell. All of the strongly responding sEC cells and PI cells are then assigned to the new CA1 cell, which from then on calculates a weighted average of their responses using a weight vector, w , which is initialized and adapted in the same way as the weight vectors of the sEC cells. The resulting place fields for some of these CA1 cells are shown in figure 7.

4. Where Have They Taken Us?

The introduction to this paper described the potential benefits of investigating computational models of place cells and their involvement in navigation. Subsequently, an overview of two recent computational models that address these issues was presented. The following sections attempt to evaluate the contributions from these models. The first section quickly describes the results from testing by the authors. The second section analyzes the differences between the two models. Then, the third section analyzes the similarities between the two models.

4.1 They Run!

Both of the models were implemented at a level of detail that allowed them to be tested in simulation or on a robot. Although this fine level of detail cannot currently be justified by the existing experimental data, it is extremely beneficial because it allows the models to be tested in ways that would not otherwise be possible. Both models successfully performed the simple tasks on which they were tested. These simple tasks serve as a good initial evaluation of the models, but little else.

Foster et al's model effectively controlled a simulated rat moving within a circular environment of 1m radius, that swims with a constant velocity of 0.3 m/s and changes direction instantaneously every 0.1s to a direction that is equal to the previous direction multiplied by 3 plus the newly commanded direction. After many trials the simulated rat using static navigation was able to learn to find a single reward within the

environment. In learning and testing it was randomly started from one of 4 locations and given no other information than the place cell representation of its current location and a reward signal at the goal location. After the metric navigation system had taken over control of the rat, the rat was able to learn to find a single reward after having found it only once previously. After this switch the static navigation system became useless.

With Arleo & Gerstner's navigation model a simulated rat successfully learned to find a single static reward within an environment using only place cell position information and reward signals at the goals. They claimed that with no further training the rat was also able to get the reward even after obstacles were placed in the environment, although this result is questionable. Their simulated rat was also able to learn to find a moved reward after significant retraining. It also learned something about how to find one reward while actually looking for another. A robot running Arleo & Gerstner's place cell model was able to learn a set of realistic looking place cells after gradually exploring an environment similar to one used in experiments on lab rats, see figure 7.

4.2 Differences Between the Two Models

In general the two models are much more similar to one another than they are to the large number of other models of hippocampal navigation that exist (Redish, 1999)(Arleo, 2000). Nonetheless some differences between the models are worth looking at.

4.2.1 Critic vs. Reward Center

Foster et al's model's use of a separate critic to create a reward signal would allow the critic to reside in a distinct module within the brain from which it could distribute the same reward signal to many different brain centers. Evidence indicates that two brain regions called the ventral tegmental area, VTA, and the substantia nigra use dopamine neurons to provide a signal that shares some similarities with δ_t to many brain areas. However, this signal seems to be strongly related to predictions of reward through time, unlike the critic, which associates a reward with the rat's location in space. For example, after training, a monkey might come to expect a reward exactly 10 seconds after seeing a visual cue (Schultz et al, 1997). Additionally, Foster et al's model would imply that these reward centers create reward maps from place cells, which would require extensive projections from the hippocampus to the VTA or substantia nigra.

Arleo & Gerstner's model only has an actor, which one would expect be closely tied to motor control rather than a center responsible for sending reward signals. However, the model did use a reward signal that would tell the actor when to forget what it had previously learned. This seemingly ad hoc addition appears to say more about the need for navigation systems to recognize when to switch strategies or the failings of static navigation, than it does about a general reward communication center.

4.2.2 The Use of Self Motion Information

Self-motion information can come from a variety of sources including inertial information, information about the static position and motion of the limbs, and knowledge of recently sent motor commands. Both models critically rely on the existence of accurate self-motion information. Foster et al's model uses self-motion information after the calculation of the place cell outputs. This information is used to learn a mapping from the place cell outputs to an (x,y) coordinate. In contrast, Arleo & Gerstner's model uses self-motion information to help produce better place cells in two important ways. First, the orientation system, which uses some self-motion information, is important to the construction and subsequent recognition of scenes by visual place cells. Second, self-motion information is used by the path integration system to

provide an additional estimate of the rat's current location. Self-motion information could be used both before and after the place cells, but new experimental results force one to reconsider the importance of self-motion information in rat navigation.

4.2.3 Path integration is not where it's at!

Strong experimental evidence shows that the output of CA1 place cells is the result of computations involving multi-modal sensory input. Anatomically, many cortical areas responsible for a variety of different modalities project to the hippocampus. Experimentally, CA1 place fields remain fairly stable even when rats navigate around a familiar environment in the dark. Most researchers including Arleo & Gerstner have sighted navigation in the dark as evidence for the importance of path integration as an input to place cells. However, true to Feynman's comments in Appendix A, a very exciting experiment published in 2000 by respected researchers in the journal *Hippocampus* gives convincing evidence that olfactory information rather than self-motion information has been used to stabilize the place fields of rats in the dark. This work used modern multi-electrode recording techniques with a carefully designed scientific experiment to contradict years of research. Wiping the floor of the environment with a moist cloth, and hence removing odors, was the key to this experiment. It seems that rats might actually deposit their own scents in an environment as they explore it (Save, 2000).

The results from this experiment indicate that a substantial part of Arleo & Gerstner's model is inaccurate. Given this evidence, as well as standard engineering knowledge that straight line displacements are very difficult to estimate using inertial information, one must also be skeptical of Foster et al's algorithm for learning a coordinate system. Since it looks as though their test may have used a noise free self-motion signal and discrete time steps it says little about the feasibility of their system. Any problems could be exacerbated by the use of long memory traces in their learning algorithm, which they chose specifically because they assumed the self-motion information would be accurate.

4.2.4 Where did the time go?

Foster et al tuned their model to fit some data describing the time it took rats to find the reward in various stages of learning, see figure 9. Their model spent all of that time learning how to move at each location when given an idealized place cell output. Real place fields take time to stabilize within a new environment and have varied shape. For example, place fields often have strange shapes at the boundaries of environments. It is unclear how their model would perform given these changes, although it is unlikely their results would still fit the experimental data.

In contrast, Arleo & Gerstner's place cell model required the robot to spend a lot of time exploring the environment in order to learn a stable place cell representation.

These are two different types of learning that could occur during experiments in which rats explore an environment and gradually improve their performance. Of course both types of learning might be occurring simultaneously, or rats might be learning something completely different. In fact, given current experimental data one could propose the unlikely theory that rats feel compelled to wander around a new environment even though they are learning nothing, perhaps to lay scents or scare away competitors.

Single episode learning within water mazes, described below, provides some weak evidence that, unlike the navigation methods proposed by Foster et al and Arleo & Gerstner, rat navigation uses exploration time to learn about the environment, not a particular reward within the environment.




Figure 9. The graph on the left shows escape latencies for rats in a water maze with a stationary submerged platform. The graph on the right shows escape latencies for rats in the same water maze with a submerged platform that moves between days. (Foster et al, 2000)

4.2.5 The Metric System

The most noticeable difference between the two navigation systems is that Foster et al's eventually uses metric based navigation while Arleo & Gerstner's does not. This distinction allows Foster et al's model to avoid conflicting with the results of recent water maze experiments. On the other hand, without suitable modifications Arleo & Gerstner's model is invalidated by these experiments. In these water maze experiments rats swim around a 2 meter diameter cylindrical tank of opaque liquid looking for a submerged 11cm diameter platform that will allow them to rest and escape. Assuming that the water is not stagnant, this experimental apparatus probably gets rid of problems with olfactory cues. As can be seen in the left graph of figure 9, rats can learn to find a stationary platform over the course of approximately three days, consisting of 12 trials total. These results can be modeled by Arleo & Gerstner's algorithm. But as shown in the right part of figure 9, rats can also learn to find a platform that moves randomly between days in about the same length of time. Given the task depicted by the graph on the right side, Arleo & Gerstner's model would fail.




Figure 10. This figure shows four successful paths by rats to the goal platform after a single example in the water maze test, where the small circle represents the goal platform. The two paths on the left are from real rats (Steele & Morris, 1999). The two paths on the right are from Foster et al's simulated rats using metric navigation. Notice the strong contrast in directness of the paths.

Additionally, as previously discussed the stability of Foster et al's learned coordinate system is questionable. Even if we ignore issues of stability due to poor self-motion information and continuous time, metric navigation still has weaknesses in its generality as a navigation strategy. Moreover, non-metric topological algorithms could effectively solve the same problem.

When instead of looking at the latencies depicted in figure 9, one examines the limited path information provided by Steele and Morris, shown in figure 10, yet another explanation appears to be feasible. The rat could simply learn the following strategy: First, when the rat discovers a platform it takes note of distal cues in the direction of the nearest wall. On the next trial the rat swims away from the wall toward the center of the pool. Once near the center of the pool the rat looks around for distal cues that match the cues it remembers from the previous trial. Finally, the rat swims toward those distal cues until it reaches the platform or fails to find the platform in the general area. If it fails to find the platform it will either go back to the center and try again or randomly search for the platform. This strategy does not require metric navigation or even a topological map, yet it could still involve the hippocampus in learning the sequence of behavior and the distal cues. This strategy could be augmented by estimation of the distance of the platform from the wall to which it is closest. The place cell model by Hartley et al (2000) discussed below lends additional credibility to this type of strategy because of its use of pre-CA1 receptive fields that are sensitive to walls. Furthermore I believe researchers have remarked at the propensity of rats to more quickly find platforms positioned in the center of the water maze (although I can't remember where I read this!).

4.3 Similarities Between the Two Models

These two models share some features that are worth looking at in isolation.

4.3.1 Simple Learning Rules

The computations performed by the models would be very easy to implement in neural hardware. The models use functions that compute simple linear combinations of their input, which could be computed by a single neuron. Likewise, the update rules for the weight vectors match well with simple neural learning that has been shown to exist in a variety of neural systems, such as hebbian learning and classical conditioning. Based on extensive studies of the biology of very simple neural systems, good evidence exists for a small library of operations used by neurons to change their responses to input (Kandel et al, 2000).

However, given science's limited knowledge of the capabilities of complex arrangements of neurons, simple computations with obvious neural implementations are not necessary for a good computational model. Science is currently unaware of most of the computations performed within the rat brain.

4.3.2 Place Cells

4.3.2.1 Radial Basis Functions

Both models use the output of place cells to represent the current position of the rat as a set of widely tuned, redundant basis functions. It is interesting that by exploiting this particular representation they are able to use very simple linear learning methods to efficiently and robustly estimate fairly complicated functions of the original input. Among other properties this representation allows them to smoothly interpolate between discrete weight values and biases their function approximations to be smooth. Although the functions they estimate seem questionable with respect to navigation, this trick might be used by the rat's brain for other tasks. This interpretation of the place cell output relates strongly to work on radial basis functions by Poggio and others, work that has more recently evolved into general kernel based learning methods (Bishop, 1995).

4.3.2.2 Abstract Neighborhoods

A clever feature of this method that is not heavily emphasized in the papers relates to experimental observations about how place cells change when obstacles to movement are introduced into an environment. Multi-electrode recordings have demonstrated that place cells very rarely cross over a wall. Moreover, place cells cease firing when a newly introduced wall crosses over their receptive field (Wilson & McNaughton, 93). Within the proposed learning algorithms the broad tuning of place cells is used to efficiently transfer learned information to neighboring locations. Due to the response of place cells to obstacles, place cells that are close together in the environment will not necessarily be close to each other in the place cell representation, thus reflecting a more abstract notion of neighborhood. Consequently the functions will not generalize their values across boundaries or major obstacles. Similarly, TD learning will not transfer values across these locations because they will not be visited in succession. Despite these differences the coordinate information should be properly transferred albeit more slowly (Dayan, 1991). One could imagine this type of system working within more abstract states spaces represented in the same way. Potentially this type of learning could use CA1 hippocampal outputs for general sequence learning, but experimental data to support this type of hypothesis is currently weak.

4.3.2.3 Generalized Snapshots

Arleo's model of place cells is significant in that it demonstrates the ease with which the current visual input can be captured and recognized from a first person perspective, so as to produce fairly well-behaved place fields. He is not the first to demonstrate this property (Fuhs et al, 1998), but at this point further experimental evidence of this kind is not wasted.

One could speculate that an intelligent creature would benefit by collecting first person perspective snapshots of a variety of sensory inputs and their combinations throughout the day. Each of these multi-modal snapshot cells would then be responsible for responding to the occurrence of sensory combinations that match its captured input. Besides simple sensory input, these snapshot cells could also take snapshots of higher-level percepts. A limited number of these snapshot cells would be necessary to sufficiently represent the important moments during a day and then could be consolidated at night into neocortical structures. Two keys to success that might make this feasible come from demonstrations such as Arleo's. First, snapshots from a first person perspective can make capture and matching relatively simple, since no search over transformations is required. Second, although having a large set of recruited snapshot detectors constantly looking for matches might appear to be computationally expensive, neural hardware makes it easy. More extreme speculation, can lead one to imagine that this type of computation might be responsible for episodic memory in humans and for representing the context of experienced situations.

These types of speculations should be used to motivate well grounded investigations into the computations that lead to the output of place cells and other CA1 pyramidal cells in the rat hippocampus. Hartley et al have produced especially nice work investigating the formation of place cells. Rather than simply resorting to currently popular vision algorithms and showing pictures that qualitatively match place fields (a blue square with a small red splotch surrounded by some yellow and green), they have used multi-electrode recordings to carefully investigate how the shape of place fields change when the environment is changed. They have been able to make models that predict how a place cells receptive field will change when the rat is placed within some types of new environments. These place cells use linear combinations of receptive fields that are sensitive to the distance and directions of walls and barriers. This is the type of work that should get attention and priority (Hartley et al, 2000).

4.3.3 Static Navigation

As described above, in its purest form the static navigation algorithm is unable to account for learning by one example. But with or without this failing, the static navigation algorithm is unappealing. This section makes some attempt to understand what is troublesome about it.

4.3.3.1 It has limited ability to generalize what it has learned.

Sequences of actions might be better represented as action sequences rather than policies controlled by the current input position. A composition operator is notably missing from the static navigation models. And the potential for linear combinations of policies as used by some systems, seems like a poor composition operator (Redish, 1999). One could imagine a sequence of actions learned in one context that would have value in other contexts, yet with the representation proposed by Foster et al and Arleo & Gerstner these subsequences would be trapped within global policy representations and tied to the states of a particular environment. These representations have limited ability to generalize what they've learned to new environments and tasks.

4.3.3.2 If it has well represented states, why can't the system just search through them?

From a classical AI perspective it is disconcerting to see such an inefficient use of a state space representation. While an agent explores an environment it would be possible for it to learn the actions that locally connect the coarsely coded states represented by the hippocampal place cells. Once the agent has a representation of the states and the actions that move between them locally, any type of classical planning method could search for sequences of actions that would transport the agent from location A to location B. Can rats imagine? Maybe rats have some type of planner that can make the hippocampus believe the rat is at a particular location and hence search through sequences of actions. Issues related to uncertainty, neural hardware, and the coding strategy used by the hippocampus might make this navigation strategy difficult.

4.3.3.3 How would this navigation method fit into the whole system of navigation?

Another point of concern is the lack of focus on how static or metric navigation would integrate with other types of navigation, such as taxon and praxic navigation. Taxon is the most primitive and widely used form of navigation in creatures. Bacteria, humans and all creatures in between use this form of navigation extensively. In an evolutionary sense and a daily use sense taxon navigation is the foundation on which other forms of navigation are built. It would seem reasonable for these other forms of navigation to take advantage of the systems that were already in place for taxon navigation. For example, strong evidence points to the superior colliculus's role in taxon navigation. The SC has registered egocentric maps of multiple modalities stacked on top of one another and topologically arranged. If a location in the SC is stimulated the animal will orient to an associated location in the immediate environment. (Redish, 1999). Perhaps these registered egocentric maps could be stimulated by an allocentric navigation system so as to create a virtual stimulus in the environment? Along these same lines a vector describing the exact direction and distance to a location in an environment is unnecessary for navigation, especially for small environments. A direction in which to travel would be sufficient for most tasks, since the rat is unable to teleport. Given the complexity of natural environments vectors representing long distances are unlikely to be very meaningful for a rat, since straight line travel will rarely be possible over long distances. Something like a logarithmic resolution of the distance would probably be sufficient, since the differences between two long Euclidean distances are insignificant relative to the same difference between two short Euclidean distances.

4.3.3.4 What ethologically relevant situations would require static navigation?

Usually it would seem that a navigation system that needs to move to hidden stimuli, would usually be operating at more of a topological level of abstraction. Once the rat gets close enough to most relevant stimuli, it will sense it and be able to use taxon navigation. Rat's have extremely sensitive sensory devices that, as researchers have discovered, are very difficult to fool.

4.3.4 Reinforcement Learning

The fundamental algorithms behind both navigation systems come from reinforcement learning. Clearly learning that is sensitive to reinforcement is essential for successful navigation, so in this sense good models will involve some form of reinforcement learning. (Sutton & Barto, 1998) As previously mentioned computations that are at least loosely related to TD learning occur in the VTA and the substantia nigra, so we might expect successful models of the future to include this system. (Schultz et al, 1997) Perhaps, as with these two models, a form of TD learning will even prove to be useful in learning simple functions based on the hippocampal cell outputs. However, static navigation as described in these models is unsatisfying and thus within this specific context it is difficult to imagine the use of TD learning. More exciting possibilities might involve using reinforcement learning to learn sequences of hippocampal activity that can be composed with one another.

5. How Might We Get There Faster?

As the previous analysis shows, combining experiments with explicit computational models can help produce better experiments and better computational models. In essence computational models with enough detail to be implemented can serve as the scientific theory that interacts with the scientific experiments. In the classic progression of science this iterative system gradually converges to a better understanding of the topic being studied.

However, if researchers were stuck using the same tools as 10 years ago this process of convergence would be excruciatingly slow. This concluding section will provide a quick overview of some promising directions in place cell research as well as some ideas for methods that might accelerate the process.

5.1 Multi-electrode Recording

As reviewed in Appendix C, Wilson and McNaughton did pioneering work on the application of multi-electrode recording to the study of rat place cells.

Many variations on multi-electrode recording methods have been developed in subsequent years. Improvements in these technologies have increased the period of successful implantation to the scale of months, lowered the complexity and cost of the apparatus, increased the number of neurons that can be monitored, and increased the number of distinct brain regions that can be monitored simultaneously. It is worth noting that the cost of computation has played a role in this progress, since setups have progressed from more than 7 computers for data acquisition for Wilson & McNaughton in 1993, to a single computer in 1997 (Nicollelis, Ghazanfar, et al, 1997).

With these improvements to the technology, the application of multi-electrode recording is becoming widespread. For example, as briefly mentioned in the introduction, the work of Wood et al (1999) has used multi-electrode arrays to find CA1 pyramidal cells that are responsive to elements of experience other than the rat's current location. As described above, the work of Hartley et al is another excellent example of the types of experiments with which this technology can assist. Chronic implantation of larger numbers of electrodes should give researchers the ability to collect much more thorough data on how hippocampal place fields change in response to a wide variety of input stimuli. Likewise, multi-site multi-electrode implantation may allow researchers to monitor cells in the entorhinal cortex at the same time that they watch CA1 pyramidal cells. This type of work should help decipher the computations leading to place

fields. Also, small wireless implementations of multi-electrode arrays could have distinctive benefits, since they could support more ethologically valid studies of place cell activity. One weakness of most research so far is the simplicity of the environments within which the rats move. More portable arrangements might encourage researchers to validate previous results in more ecologically valid environments.

5.2 Quantifying the Input

Once again, as emphasized by Feynman's comments in Appendix A. and Save et al's results (2000) on olfaction, knowing what information is used as input to the computations is essential. The ideal situation would be to create virtual reality systems for rats, like the systems used in experiments with humans (Maguire et al, 1998). In order to study place cells, rats might actually need to move, in which case rat VR would be less effective. Clearly, when the rat must move, careful procedures must be used to remove, or model, the influence of olfaction, perhaps by severing the nerves to olfactory processing centers in the rat brain. Additionally other sensory information must be carefully controlled. Besides rat VR, another option that might give researchers more control over the visual input to the rat, would be a video projection system that projects visual stimuli onto a dome surrounding the rat's environment. Another option would be to collect images from all places within the environment by an automated device that traverses the test environment. A researcher could collect a database of images, or create a high quality 3D model, that fully specifies the view a rat would have from any position within the test environment. As long as the content of the room and the lighting is carefully controlled, recordings of the rat's head position, body position, and eye position over time would allow researchers to create a virtual rat that would facilitate studies of rat navigation and place cell responses to visual stimuli. These databases could then be shared over the web by researcher's producing computational models as well as experimentalists. Besides Hartley et al's research, an example of the importance of this detailed data is evident in the lack of data provided by Steele and Morris (1999). Although they collected full path information by an overhead camera for the rats in the water maze they only showed two useful example paths in their paper, which was insufficient to evaluate the proposed strategy for navigation. Detailed information can help eliminate hypotheses.

5.3 Models of Microstructure




Figure 11. 3D reconstructions of dendrites within a small volume of hippocampal tissue from CA1. The diameter of each of these dendrites is smaller than a micron. (Fiala & Harris, 2001)

Unlike the two models reviewed here, some computational theories do attempt to model the hippocampus at a level closer to the microstructure. Some work produces detailed compartmental models that simulate the electrical properties of hippocampal structures using large scale neural simulation packages like Genesis (Menschik, 1999). Other work, is more like that of Ascoli et al (1999), which attempts to use databases of 3d cell reconstructions to make statistical models of different hippocampal cell morphologies, from which large scale neural assemblies can be generated which are qualitatively similar to the hippocampus.

Unfortunately, even though the hippocampus is one of the most studied brain structures, current anatomical data appears to inadequately support such efforts. This inadequacy is related to both the previous lack of low cost methods for working with large amounts of detailed 3D data and the labor-intensive processes currently required to create the 3D reconstructions. This situation should improve as computer driven microscopy becomes more wide spread and refined. For example, serial electron microscopy has been used to create 3D reconstructions of neurons within small volumes of hippocampal tissue from CA1 at remarkable resolutions, but using labor-intensive processes, see figure 11 (Fiala & Harris, 2001). As robotics and 3D image processing automate these processes, significantly better data should become available.

5.4 Genetics

Although this summary cannot include every promising avenue, genetic techniques deserve to be mentioned. Genetic engineering has allowed researchers to precisely target rat hippocampal cells and alter their properties through their genes. For example, many interesting results relating to the modification of NMDA sensitive receptors, which are important for some types of long term learning, have been published. (Kandel et al, 2000)

6. Conclusion

For a variety of reasons the rat hippocampus in general and rat place cells in particular are worth studying. Place cells provide a hook into understanding very important aspects of the computations performed by brains. Explicit computational models are an essential part of this study, along with careful scientific practices, and the creation of new methods for observing the input, output, and processes of the rat brain. The future of place cell research will be exciting to watch ...don't forget.

Appendix A

An apt quote from Richard Feynman:

“... there have been many experiments running rats through all kinds of mazes, and so on -- with little clear result. But in 1937 a man named Young did a very interesting one. He had a long corridor with doors all along one side where the rats came in, and doors along the other side where the food was. He wanted to see if he could train the rats to go in at the third door down from wherever he started them off. No. The rats went immediately to the door where the food had been the time before.

The question was, how did the rats know, because the corridor was so beautifully built and so uniform, that this was the same door as before? Obviously there was something about the door that was different from the other doors. So he painted the doors very carefully, arranging the textures on the faces of the doors exactly the same. Still the rats could tell. Then he thought maybe the rats were smelling the food, so he used chemicals to change the smell after each run. Still the rats could tell. Then he realized the rats might be able to tell by seeing the lights and the arrangement in the laboratory like any commonsense person. So he covered the corridor, and still the rats could tell.

He finally found that they could tell by the way the floor sounded when they ran over it. And he could only fix that by putting his corridor in sand. So he covered one after another of all possible clues and finally was able to fool the rats so that they had to learn to go in the third door. If he relaxed any of his conditions, the rats could tell.

Now, from a scientific standpoint, that is an A-number-one experiment. That is the experiment that makes rat-running experiments sensible, because it uncovers that clues that the rat is really using - - not what you think it's using. And that is the experiment that tells exactly what conditions you have to use in order to be careful and control everything in an experiment with rat-running.

I looked up the subsequent history of this research. The next experiment, and the one after that, never referred to Mr. Young. They never used any of his criteria of putting the corridor on sand, or being very careful. They just went right on running the rats in the same old way, and paid no attention to the great discoveries of Mr. Young, and his papers are not referred to, because he didn't discover anything about the rats. In fact, he discovered all the things you have to do to discover something about rats. But not paying attention to experiments like that is a characteristic example of cargo cult science.”

-Adapted from a Caltech commencement address given in 1974 from the book from "Surely You're Joking, Mr. Feynman!" ISBN 0393316041

Appendix B

A Review of Wilson and McNaughton (1993)

In 1993 Wilson and McNaughton published results from experiments that used multi-electrode techniques to simultaneously record the activity of large numbers of neurons within the pyramidal layer of CA1 in rats. These results were broadly significant in that they helped pioneer the use of multi-electrode recording techniques for the study of neural structures in vivo. Since this research multi-electrode recording has become a widely used tool for neuroscience that has shown significant benefits over single electrode techniques, which have been used since 1926. Contributions were also made toward understanding place cells, since the work gave both new insight into the concurrent time dependent interactions of place cells and strong corroborating evidence for previously stated characteristics of place cells. (Nicolelis, 1997)

One of the more influential results of this work was the accurate reconstruction of the rat's location within an environment based on the recorded firing rates of the population of neurons. Specifically, for each of three rats they recorded the current position of the head at 20Hz within a 124cm by 62cm rectangular box, along with the action potentials of over 70 distinct neurons digitized at 33kHz. They then estimated a function that would output the rat's location when given information about the firing rates of the neuron population.

The primary weakness with their approach to this estimation is that they tested their estimated function using the same data on which they trained the function. This puts the burden of evidence on the type of function they used, since if they were to simply use a look-up-table containing all the data they would meaninglessly get nearly perfect reconstruction.

For each location, they calculated the mean firing rate of a set of recorded neurons over the entire episode, where the mean firing rate over a period of time is just the total number of action potentials recorded during that time divided by the length of that period. They do not state how they quantized the environment to produce the set of locations they used, although the number could be as low as 6x12 given their claim of 5cm intrinsic tracking error. They then computed the mean firing rate of this same set of neurons over brief windows of time. When given the mean firing rate over a window of time as input, the function would return the location with the closest overall mean firing rate.

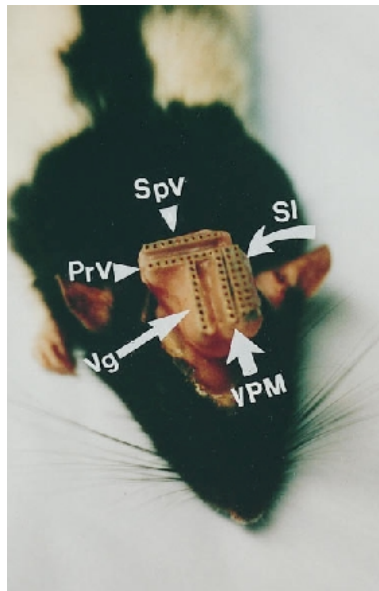
Given the issues with their approach to reconstruction, a much more convincing form of evidence for the role of place cells has been based on histograms of the firing of individual neurons as a function of location. These visualizations of the data very clearly communicate several of the properties of place cells. They are sparse, non-topologically arranged, well localized and unimodal.

More recent work has performed more elaborate estimations of the locations of rats based on place cell populations (Zhang et al, 1998)(Brown et al, 1998). The broader success of multi-electrode recording and methods of population decoding is visible throughout the neural sciences. For example, rats and monkeys have been able to control robotic devices in real-time by thought alone. Due to the constraints on causal, real-time control this work very clearly demonstrates the successful decoding of motor control information from a population of neurons (Chapin et al, 1999)(Wessberg et al, 2000).

A quick overview of the multi-electrode technology and methodology:

For each rat they implanted 12 tetrodes, where 4 insulated wires were twisted together to make a tetrode with a total diameter of 40 microns. For reference, a pyramidal neuron is typically ... The tetrode configuration is important because by monitoring the relative potential differences among the tips of the four wires, they were able to detect and isolate the action potentials of neurons within a small volume of tissue. Through this method each tetrode was able to resolve the activity of 5 to 20 single neurons in close proximity to the tetrode. Consequently, with one of the rats they were able to monitor 148 neurons with 48 wires. (Wilson & McNaughton, 1993) (O'Keefe & Burgess, 1996) (O'Keefe, 1999)

In order to successfully record the activity of CA1 pyramidal neurons as opposed to other types of neurons in the CA1 region or neurons from another region altogether, several techniques are used. First, tools and techniques have been developed that allow the researcher to secure the rat's skull and place the brain into a common coordinate system within which the coordinates of most major brain locations have been documented. By using tools precisely aligned to this coordinate system the researcher is able to move the electrodes into the desired general area. Due to the hippocampus's large sheet like structure that runs roughly parallel to the contour of the head, depth is the only axis with considerable uncertainty. Second, the researchers monitor the output of the electrodes as they fully adjust their positions and look for the characteristic firing of CA1 pyramidal neurons, which are referred to as complex-spike cells due to their intricate firing properties. As O'Keefe makes clear, it is especially important for researchers to distinguish between pyramidal cells and theta interneurons. Third, after the experiments the animals are euthenized and the electrode locations are confirmed. (O'Keefe, 1999) (Nicolelis, 1997)



a rat with an implanted multi-electrode
(Nicolelis, Ghazanfar et al, 1997)

Bibliography

Araujo IET, Rolls ET, Stringer SM. A View Model Which Accounts For the Spatial Fields of Hippocampal Primate Spatial View Cells and Rat Place Cells. *Hippocampus* 2001;11:699-706.

Arleo A. Spatial Learning and Navigation in Neuro-Mimetic Systems, Modeling the Rat Hippocampus. Ph.D. thesis, Department of Computer Science, Swiss Federal Institute of Technology Lausanne, EPFL, Switzerland, 2000.

Arleo A, Gerstner W. Spatial Cognition and Neuro-Mimetic Navigation: A Model of Hippocampal Place Cell Activity. *Biological Cybernetics* 2000;83:287-299.

Ascoli GA. and perspectives in computational neuroanatomy. *Anat Rec.* 1999 Dec 15;257(6):195-207.

Bishop CM. *Neural Networks for Pattern Recognition*. Oxford: Oxford University Press; 1995.

Brown EN, Frank LM, Tang D, Quirk MC, Wilson MA. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *J Neurosci.* 1998 Sep 15;18(18):7411-25.

Bulman-Fleming, B. Department of Psychology, University of Waterloo (Waterloo, Ontario). 2000. Psychology 261, Lecture 15 notes. World Wide Web URL: <http://www.arts.uwaterloo.ca/~bfleming/psych261/lec15no7.htm>.

Bunsey M, Eichenbaum H. Conservation of hippocampal memory function in rats and humans. *Nature* 1996;379:255-257.

Burgess N, Donnett JG, Jeffery KJ, O'Keefe J. Robotic and neuronal simulation of the hippocampus and rat navigation. *Phil. Trans. Roy. Soc. B* 1997;352:1535-1543.

Burgess N, Jackson A, Hartley T, O'Keefe J. Predictions derived from modelling the hippocampal role in navigation. *Biological Cybernetics* 2000;83:301-312.

Burgess N, Maguire EA, Spiers HJ, O'Keefe J. A temporoparietal and prefrontal network for retrieving the spatial context of lifelike events. *NeuroImage* 2001;14:439-453.

Chapin, J.K, Moxon, K.A., Markowitz, R.S. and Nicolelis, M.A.L. (1999) Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nature Neurosci.*, 2:664-670.

Chudler, EH. Department of Anesthesiology, University of Washington (Seattle, WA). 2001. Brain Facts and Figures. World Wide Web URL: <http://faculty.washington.edu/chudler/facts.html>

- Dayan P. Navigating through temporal difference. In: Lippmann RP, et al, editors. In: Lippman RP, Moody JE, Touretzky DS, editors. *Advances in Neural Information Processing System 3*. City: Publisher; 1991. p. 464–470.
- Edgar PF, Douglas JE, Knight C, Cooper GJ, Faull RL, Kydd R. Proteome map of the human hippocampus. *Hippocampus* 1999;9(6):644-50.
- Eichenbaum H. Is the rat hippocampus just for “place”? *Curr Opin Neurobiol* 1996;6:187–195.
- Fiala JC, Harris KM. Extending unbiased stereology of brain ultrastructure to three-dimensional volumes. *Journal of the American Medical Informatics Association*. 2001;8(1):1-16.
- Fiala JC, Spacek J, Harris KM. Laboratory of Synapse Structure and Function at Boston University (Boston, MA). 2001. Hippocampus of the Rat. World Wide Web URL: <http://synapses.bu.edu/anatomy/hippo/hippo.stm>.
- Foster DJ, Morris RGM, Dayan P. Models of hippocampally dependent navigation using the temporal difference learning rule. *Hippocampus* 2000;10:1-16.
- Fuhs MC, Redish AD, Touretzky DS. A visually driven hippocampal place cell model. In J. Bower (Ed.), *Computational Neuroscience: Trends in Research*, New York: Plenum Publishing; 1998; 379-384.
- Georgopoulos A, Schwartz A, Kettner R. Neuronal population coding of movement direction. *Science* 1986;233: 1416-1419
- Hartley T, Burgess N, Lever C, Cacucci F, O'Keefe J. Modeling place fields in terms of the cortical inputs to the hippocampus, *Hippocampus* 2000;10:369-379.
- Huxter JR, Thorpe CM, Martin GM, Harley CW. Spatial problem solving and hippocampal place cell firing in rats: control by an internal sense of direction carried across environments. *Behav Brain Res*. 2001 Aug 27;123(1):37-48.
- Kandel ER, Schwartz JH, Jessell TM. *Principles of Neural Science*. McGraw-Hill; 2000.
- Louie K, Wilson MA. Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron*. 2001 Jan;29(1):145-56.
- Maguire EA, Burgess N, Donnett JG, Frith CD, Frackowiak RNA, O'Keefe J. Knowing where and getting there: a human navigation network. *Science* 1998;280:921-924.
- Menschik ED. Cholinergic Neuromodulation in Hippocampal Function and Disease: From single-cell biophysics to network simulations. PhD Thesis, University of Pennsylvania, Neuroscience Graduate Group, April 1999.
- Nicolelis MA, Fanselow EE, Ghazanfar AA. Hebb's dream: the resurgence of cell assemblies. *Neuron*. 1997 Aug;19(2):219-21.
- Nicolelis MA, Ghazanfar AA, Faggin BM, Votaw S, Oliveira LM. Reconstructing the engram: simultaneous, multisite, many single neuron recordings. *Neuron*. 1997 Apr;18(4):529-37.

O'Keefe, J. Do hippocampal pyramidal cells signal non-spatial as well as spatial information? *Hippocampus* 1999;9:352-364.

O'Keefe J, Burgess N. Geometrical determinants of the place fields of hippocampal neurons. *Nature* 1996;381:425-428.

O'Keefe J, Burgess N. Theta Activity, Virtual Navigation and the Human Hippocampus. *Trends In Cognitive Science* 1999;3:403-406.

O'Keefe J, Burgess N, Donnett JG, Maguire EA, Jeffery KJ. Place cells, navigational accuracy, and the human hippocampus. *Philosophical Transactions of the Royal Society B* 1998;353,1333-1340.

Pouget A, Dayan P, Zemel R. Information processing with population codes. *Nat Rev Neurosci.* 2000 Nov;1(2):125-32.

Redish DA. *Beyond the Cognitive Map*. Cambridge: MIT Press; 1999.

Salinas E, Abbott L. Vector reconstruction from firing rates. *J Comp Sci* 1994;1:89-107.

Save E, Nerad L, Poucet B. Contribution of multiple sensory information to place field stability in hippocampal place cells. *Hippocampus*, 2000;10(1):64-76.

Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science*. 1997 Mar 14;275(5306):1593-9.

Spiers H, Burgess N, Maguire EA, Vargha-Khadem F, O'Keefe J. Bilateral hippocampal pathology impairs topographical and episodic memory but not visual pattern matching. *Hippocampus* 2001.

Spiers HJ, Burgess N, Maguire EA, Baxendale SA, Hartley T, Thompson P, O'Keefe J. Unilateral temporal lobectomy patients show lateralised topographical and episodic memory deficits in a virtual town. *Brain* 2001.

Steele RJ, Morris RG. Delay-dependent impairment of a matching-to-place task with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5. *Hippocampus*. 1999;9(2):118-36.

Sutton RS, Barto AG. *Reinforcement Learning*. Cambridge: MIT Press; 1998.

Wessberg J, Stambaugh CR, Kralik JD, Beck PD, Laubach M, Chapin JK, Kim J, Biggs SJ, Srinivasn MA, Nicolelis MAL. Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 2000;408:361-365.

Wilson MA, McNaughton BL. Dynamics of the hippocampal ensemble code for space. *Science* 1993;261:1055-1058.

Wilson MA, McNaughton BL. Reactivation of hippocampal ensemble memories during sleep. *Science* 1994;65:676-679.

Wood ER, Dudchenko PA, Eichenbaum H. The global record of memory in hippocampal neuronal activity. *Nature* 1999;397:613–616.

Zhang K, Ginzburg I, McNaughton BL, Sejnowski TJ. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J Neurophysiol.* 1998 Feb;79(2):1017-44.