

Modeling Place Fields in Terms of the Cortical Inputs to the Hippocampus

Tom Hartley,^{1,2*} N. Burgess,^{1,2} C. Lever,²
F. Cacucci,² and J. O'Keefe^{1,2}

¹*Institute of Cognitive Neuroscience, University College London, London, UK*

²*Department of Anatomy and Developmental Biology, University College London, London, UK*

ABSTRACT: A model of place-cell firing is presented that makes quantitative predictions about specific place cells' spatial receptive fields following changes to the rat's environment. A place cell's firing rate is modeled as a function of the rat's location by the thresholded sum of the firing rates of a number of putative cortical inputs. These inputs are tuned to respond whenever an environmental boundary is at a particular distance and allocentric direction from the rat. The initial behavior of a place cell in any environment is simply determined by its set of inputs and its threshold; learning is not necessary. The model is shown to produce a good fit to the firing of individual place cells, and populations of place cells across environments of differing shape. The cells' behavior can be predicted for novel environments of arbitrary size and shape, or for manipulations such as introducing a barrier. The model can be extended to make behavioral predictions regarding spatial memory. *Hippocampus* 2000;10:369–379. © 2000 Wiley-Liss, Inc.

KEY WORDS: model; space; rat; neural network

INTRODUCTION

Place cells in the hippocampus of the rat exhibit location-specific firing. A given cell fires only when the animal is in a particular part of its environment (the place field). When a place cell is recorded in a series of rectangular environments, the location of peak response shows some regularity across environments, e.g., often maintaining a fixed distance from the two nearest walls (O'Keefe and Burgess, 1996). Regularity in the locations of place fields has also been observed across a variety of different-shaped environments (e.g., squares, circles; Lever et al., 1999). Although this regularity appears to be an important and reliable feature of place cells in these experiments, other experiments have shown "remapping," i.e., unrelated patterns of firing in different environments (Muller and Kubie, 1987; Quirk et al., 1992) or mixtures of regularity and remapping (Skaggs and McNaughton, 1998). The occurrence of shape-based remapping may well depend on extensive

experience in both environments (Lever et al., 1999 and unpublished data). Our aim is to explain as much as possible of the regularity in the initial pattern of place-cell firing in new environments. Does a place cell need to learn its response to a new environment? Or can its behavior be understood as a hardwired response to the geometry of the environment? We will argue for the latter position, and describe a model of the cortical inputs required for such a response.

In the rectangular environments studied by O'Keefe and Burgess (1996), place fields can be modeled as the thresholded sum of two or more putative inputs of specific functional form. Inputs have a Gaussian tuning curve to the distance to a wall in a given allocentric direction (e.g., north, south). Each of the inputs is tuned to respond maximally when there is a wall at a specific distance along a specific direction. The breadth of the tuning increases with the distance to which it is tuned. This simple model can account for the consistency in the location of firing between environments of different size and aspect as well as various features of the shape of the fields, such as their elongation in rectangular environments compared with square (O'Keefe and Burgess, 1996; Burgess and O'Keefe, 1996). In this study it was convenient to model place fields using four inputs, each tuned to respond to distances in the directions perpendicular to each of the walls, since these were the four directions along which the shape of the environment was varied. However, in general we would not expect the directions to which inputs to a given place cell are tuned to be determined by a particular environment. Although a place cell will tend to fire when the various walls match the tuned distances of its inputs, those distances can be measured in many directions. The general model simply suggests that place cells are driven by inputs sensitive to the distance and allocentric direction of boundaries ("boundary vectors") in the environment, with several such inputs combining to produce location specificity. This model builds on previous models of place cell firing,

Grant sponsor: UK Medical Research Council.

*Correspondence to: Tom Hartley, Institute of Cognitive Neuroscience, University College London, Gower St., London WC1E 6BT, UK.

E-mail: t.hartley@ucl.ac.uk

Accepted for publication 1 May 2000

many of which involve distances and/or bearings to landmarks (e.g., Zipser, 1985; Sharp, 1991; O'Keefe, 1991; McNaughton et al., 1994; Burgess et al., 1994; Touretzky and Redish, 1996). However, these previous models do not give any special status to boundaries of the environment as compared with other landmarks and do not attempt to find the precise functional form of the inputs required to predict its behavior quantitatively.

What constitutes a boundary is a question that requires further empirical investigation, but it is clear that walls which impede movement are particularly potent cues in determining the location of peak response for most place cells (O'Keefe and Burgess, 1996). Place fields near to a small, open box forming part of a larger apparatus and moved in the laboratory frame also maintain fixed locations relative to the walls of the box (Gothard et al., 1996a,b). Other studies show that the presence of a transparent barrier to movement is sufficient to affect place-cell firing patterns (Muller and Kubie, 1987). Place fields also maintain their positions relative to the edges of a raised holding platform (O'Keefe, 1979) as it is moved in the laboratory frame, suggesting that a drop which impedes movement also constitutes a boundary. Other data suggest that as well as impeding movement, barriers must be extended in space to affect place field location. For instance, Cressant et al. (1997) showed that small objects placed inside a cylindrical environment did not affect the location of firing when they were separated, but did when they were arranged in a row. We do not assume that information concerning the location of boundaries in the environment comes exclusively from the visual modality. Although when visual inputs are available they tend to dominate other sources of information (e.g., Jeffery et al., 1997), it has long been thought that place cells will use any available external or internal information to maintain location specific firing (e.g., O'Keefe and Nadel, 1978; Hill and Best, 1981).

greatest to a boundary at distance d_i

MODEL

Framework

We model the geometric inputs to the place cells as a population of *boundary vector cells* (BVCs), each of which responds maximally when a boundary is at a particular distance and allocentric direction to the rat, as shown in Figure 1.

The receptive field of each BVC is a product of two Gaussians, one a function of distance, the other of allocentric direction. The distance tuning of the cells is narrow for cells which have a peak response to boundaries near the rat, but wider for greater distances (see Fig. 2). This is consistent with the rat being able to judge short distances more accurately than long distances (cf. Weber's law), and it means that inputs tuned to shorter distances exert more influence on place cell firing than BVCs tuned to long distances, everything else being equal. The angular extent (σ_{ang}) of each field is assumed to be constant. Thus, for a BVC i whose response is

$$\Delta f_i \propto g(r, \theta) \times \Delta\theta \quad (2)$$

to the firing rate, f_i

Field Shape and Location

Before simulating any specific cell, we can make a number of general predictions regarding place field shape and location in different environments. These are predictions that follow from the form of input representation we have postulated, and do not depend on the parameters of the model, or the specific inputs that activate a place cell.

Where the directional reference frame is anchored by the presence of salient extramaze orientation cues, the location of the field in different environments, as measured by the location of peak firing, will generally be fixed with respect to the walls that bound the environment along one or more of its sides (e.g., to the north or west).

The shape of the boundaries will affect the shape of the fields, especially near the edges of the environment. Figure 3 gives examples of the different shaped fields predicted for different environments for a cell with the same BVC inputs. Fields tend to be crescent-shaped where they are close to the perimeter of a circular environment (cf. Muller et al., 1987). The shape of fields close to barriers is also influenced by the barrier's orientation relative to the directional reference frame. For example, rotation of a square environment can produce a range of shapes from bar-shaped fields parallel to the barrier (see the "diamond" in Fig. 3) to more confined elliptical fields (see the small square in Fig. 3).

Physiological Predictions

To test the model, we make detailed physiological predictions about the behavior of a particular place cell in a novel environment. To do this, we need to solve the inverse problem: determining which BVCs drive the place cell from the firing rate map obtained in a first set of environments, and using this knowledge to predict the firing-rate maps in novel environments. Although each place cell may receive inputs from many geometry-sensitive inputs, models involving few BVCs are more likely to show robust generalization to novel environments than models involving many BVCs, as the extra degrees of freedom are likely to cause the model to overfit the data. Our aim therefore is to find a small number of BVCs that can account for the data, and thus produce robust physiological (firing rate) predictions, rather than a detailed description of the inputs at the physiological level.

For any single environment, the problem is insoluble, since there are many combinations of BVCs that could account for the observed firing. However, some BVCs (e.g., those that themselves fire strongly at the location where the peak response of the place cell is seen) are more likely to be involved in driving the place cell than others (e.g., those that fire strongly where the place cell does not fire at all), so it is possible to systematically reduce the set of BVCs considered. By fitting fields recorded in a number of environments with different boundary configurations, the number of fitting combinations can be reduced to a level where robust predictions from the model are possible.

Fitting Procedure

Firing rate maps recorded from the same place cell in a variety of environments are smoothed and normalized ($F(\mathbf{x})$). Simulated firing rate maps can then be generated for any population of BVCs (covering the range of distances and directions the rat will encounter in the experimental environments) in each environment ($f_i(\mathbf{x})$), using Eq. (2). Simulated place-cell firing maps (thresholded sums of the BVC firing rate maps) can then be fitted with the experimental data.

To determine the set of BVCs that best explain the observed data, we need to invert Eq. (3). One solution is to estimate the place field as a simple linear combination of inputs, i.e.,

$$F(\mathbf{x}) \approx \sum w_i f_i(\mathbf{x}) \quad (4)$$

This allows a direct calculation of the solution with the best fit (in the "least squares" sense) to the experimental data by inverting Eq. (4) to find the values of w_i . However, modeling place fields that can have zero firing rates in entire environments requires the nonlinear threshold of Eq. (3) rather than a linear model. We are also reluctant to abandon the simplicity of Eq. (3) by introducing the large set of variable weights (w_i in Eq. (4)). Although models obtained by this approach can fit the data well (see Burgess et al., 2000), it is necessary to model a large number of BVCs, representing an entire population of inputs which can encode for all distances and directions. This can lead to problems of overfitting and poor generalization.

Ideally, we could use a Bayesian framework to find the probability of all sets of BVC inputs, given the data and our assumptions concerning the distribution of BVCs available. This could be used to generate a probability distribution of the firing rate at each point in a novel environment. However, this does not seem readily tractable, and remains for future work.

Instead, the approach we use here is to choose a smaller population of BVCs whose firing fields ($f_i(\mathbf{x})$) are compatible with the observed data, and from them select the best-fitting combination of BVCs by "brute force" (i.e., by calculating all potential solutions). Compatible BVCs are those which represent vectors from each observed location of peak firing to the walls of the environment in all directions. Peaks were defined as contiguous regions where firing rate exceeded half the maximum rate, and the distance and direction of the peak from the walls were measured with respect to the pixel showing the maximum firing rate within each such region. As fitting each combination of BVCs is fairly computationally intensive, the number of BVCs involved was constrained to make a brute force solution practical. This was done in two ways:

1. Boundary vectors to be considered were quantized to a polar grid (with angular resolution 18° and radial resolution at four pixels or 9.76 cm).
2. Only combinations where all BVCs represented vectors orthogonal or opposed to one another were considered (i.e., there were only 2–4 BVCs in each combination). This latter constraint is reasonable, since cells representing very similar boundary vectors are likely to be redundant (their firing fields covary strongly with

one another), and BVCs with very similar bearings combine to produce extremely elongated fields, in contrast to the more symmetrical fields that are observed experimentally.

All combinations meeting the above constraints were considered. For each combination, the sum of the BVC firing fields was calculated, and the threshold T , and the global weighting parameter A (see Eq. (3)), were varied until the fit with the data was maximized (simplex minimization). Figure 4 shows the best-fitting solution obtained for each of the 28 fields described by O'Keefe and Burgess (1996), using a maximum of four BVCs. Even in the case of a model using four BVCs, only seven independent parameters are used: the four BVC distance tunings, and one direction defining the four (orthogonal) BVC direction tunings and parameters A and T .

Despite the strong constraints we applied to the combinations of boundary vector cells considered, the model fits the data remarkably well, accounting for the shape and size as well as the location of most fields.

The main discrepancies with the data were where the model predicted either firing in an environment where none was observed, or less frequently, the absence of firing where in fact firing was observed. These discrepancies suggest that place cells' firing thresholds sometimes vary from one trial to another, due to factors other than the geometry of the environment.

Having arrived at a solution to the inverse problem, the BVCs implicated in driving the place cell in one set of environments can be used to calculate the expected firing rate in a novel environment or set of environments. Figure 5 shows how the model fits a place cell's firing in four environments and how the BVCs found are used to predict its firing in four novel environments. Figure 6 shows how the model predicts accurately not just the location but also the shape of fields in a variety of environments. In particular, the model predicts that the cell, which fires close to the south wall in a rectangular environment, will produce an additional firing field in response to an east-west barrier introduced into the middle of the box. The model also correctly predicts the location of the field observed in right-angled triangles of different orientations. The firing field is extended along the hypotenuse of the triangle when the right angle is in the northeast corner, and adjacent to the south wall when the right angle is in the southeast.

Statistical and Behavioral Predictions

We can make some predictions concerning the way in which the form of representation we have suggested for the inputs to the hippocampus affects the statistical properties of a population of place cells, and the consequences that these statistical properties might have for behavior. In order to develop these ideas further, it is necessary to model a population of BVCs and a representative population of place cells with inputs from these BVCs. Clearly both entail additional assumptions which are discussed briefly below before we turn to the results of our simulations.

As boundary vector cells are purely theoretical at present, it is desirable that our predictions do not rest too heavily on our assumptions about the detailed shape and size of their receptive fields

(which are determined by the parameters σ_{an_g} , β , σ_0 , or their number and distribution in space. Nonetheless, in order to simulate a population of BVCs, we are obliged to make some assumptions about these issues, and where we believe these have important qualitative effects on our predictions, we have noted the fact. For practical reasons, we have simulated a small number of BVCs (200), which we distributed so as to cover space evenly in 20 directions, and 10 distances covering the range encountered in typical experimental environments. The parameters governing the width of the receptive fields of BVCs were the same as those used above in Physiological Predictions, and were chosen so that the overlap between adjacent fields was sufficient to allow walls at any distance or direction to activate some of the BVCs (see Appendix for more details).

The proportion of available inputs activating any given place cell is another area where we are obliged to make additional assumptions in order to simulate a representative population of place cells. In Physiological Predictions (see above), we simulated place cells with four or fewer BVC inputs to avoid overfitting the data, thus allowing us to make robust predictions. These inputs were also constrained to be separated by 90° to one another. Since we wish to remove these constraints in the following simulations, which rely on generating a representative sample of place cells with a random selection of BVC inputs, we assume that each place cell has 10 BVC inputs (5% of the total number of simulated inputs). This figure is sufficient to avoid a large proportion of elongated fields (a smaller sample of randomly chosen BVCs is also likely to sample the different directions unevenly, which produces elongated fields), while producing a good proportion of strongly localized, unimodal fields in line with experimental observation.

Statistical Properties of Populations of Place Cells

We investigated whether the statistical properties of populations of place-cell firing could be explained on the basis that the inputs to each cell are simply a random selection from the set of all possible inputs. We took the population of 28 fields recorded by O'Keefe and Burgess (1996) as our experimental data, and compared properties of their firing-rate maps with 100 simulated place cells, whose firing rate maps we calculated in each of the four rectangular environments. Each simulated place cell had inputs from 10 randomly selected BVCs (as described above). All 100 simulated cells used the same threshold value, T , chosen so that about 60% of cells, which fired in any of the boxes, would fire in all four (in line with the experimental data). To simulate real firing rates, the parameter A was set such that the median peak firing rate across all cells and environments corresponded with that observed experimentally (4.85 Hz). Simulated cells were considered not to fire where their peak firing rate in an environment was less than 1 Hz (as with the experimental data in O'Keefe and Burgess, 1996).

After thresholding the summed inputs to each simulated cell, some cells did not fire in any of the four boxes (16/100). The existence of a proportion of pyramidal cells which are silent in any particular experimental environment is consistent with experimen-



FIGURE 4. A: Data from 28 place fields reported by O'Keefe and Burgess (1996), shown for comparison with simulated cell firing rate maps shown in B. Shading indicates rate of firing, which has been normalized within each environment, so that the peak firing rate corresponds to the darkest shading. As in O'Keefe and Burgess (1996), fields with peak firing rates less than 1 Hz are not shown, and the cell was treated as if it had zero firing rate in these environments.

Pretraining of the rat occurred in the vertical rectangle, except for cells shown in the rightmost column, in which pretraining occurred in the horizontal rectangle. B: Firing rate maps of 28 simulated place cells, with 2–4 BVC inputs chosen to fit the corresponding experimental data shown in A. These plots show the actual fitted values for $F(x)$, using the same shading scale as in A.




FIGURE 4. (Continued)

tal data showing that a much larger number of cells is active under

Behavioral Predictions

The location-specificity of place cells suggests that they may play a role in guiding behavior when it is directed to a particular location (O'Keefe and Nadel, 1978). For instance, memory for the location of an object might include a representation of place based on the firing rates of hippocampal place cells. Here we need to consider the neural representation of place not in terms of the firing of single place cells, but in terms of the pattern of firing across place cells, since it is unlikely that any single cell controls behavior.

A very simple model of object location memory would be that an animal encountering an object at a particular location stores the pattern of place-cell firing rates at that time. To search for the object, the animal would focus on the location which produce the same or most similar patterns of firing in the place cells (see the "simple model" described in Burgess and O'Keefe, 1996). As shown above, the current model permits us to simulate a population of place cells, and we can thus obtain a measure of the similarity of any two locations in terms of their neural representations. We can thus use the model to make tentative predictions at the level of behavior (tentative because the relationship between neural representation and behavior may be more complex than the simple one outlined here, e.g., as in the full model of Burgess et al., 1997).

does not appear to be any trend relating size of environment to the number of cells which fire in it. This discrepancy suggests that either the input to the place cells is increased in smaller environments (i.e., inputs to the hippocampus are statistically skewed towards short distances, and the modeled population of BVCs had too few cells tuned to short distances compared to cells tuned to long distances), or the scale of the environment affects the threshold of the place cells (perhaps through interactions between place cells). The use of a single fixed threshold parameter for all cells in all environments was a useful simplifying assumption, but one which is probably excessively conservative. While the threshold remains the same, whatever the size of the environment, in larger environments there are always more inputs available to drive place cells, and hence the firing rate tends to be higher.

The model explains a large number of the statistical properties of place fields, and while the quantitative fit is in no way spectacular, it is remarkable that so many properties of place cells can be accounted for with such a simple model. Random connectivity predicts many of the variations in firing which are sometimes taken as evidence of "remapping" (e.g., a given cell fires in one environment but not in another, or produces a two-peaked field in one environment and a one-peaked field in another).

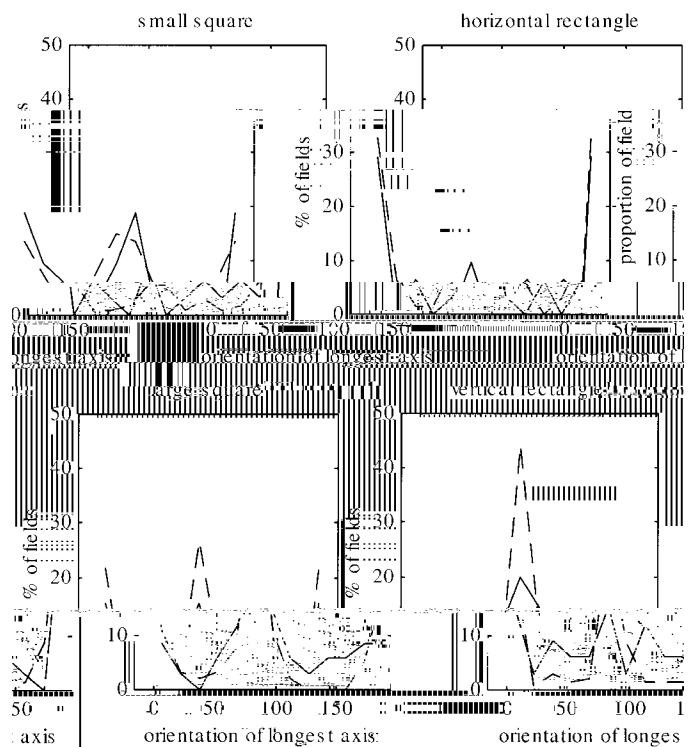


FIGURE 8. Orientation of place fields in real (solid lines; data from O'Keefe and Burgess, 1996) and simulated place fields (dashed



[REDACTED]

[REDACTED]