

Commentary

Cognitive Graphs, Resistive Grids, and the Hippocampal Representation of Space

NEIL BURGESS and JOHN O'KEEFE

From the Department of Anatomy, University College London, London WC1E 6BT, United Kingdom

The hippocampal formation appears to play an important role in the ability of rats and other mammals to perform sophisticated spatial tasks, beyond simply following a familiar route or approaching a visible cue. Evidence for this role came from the observation of selective spatial deficits after hippocampal lesions (O'Keefe et al., 1975; O'Keefe and Nadel, 1978; Morris et al., 1982; Sutherland et al., 1982) and from the spatially selective firing of hippocampal complex spike cells (O'Keefe and Dostrovsky, 1971; McNaughton et al., 1983; Muller et al., 1987) ("place cells"), and led to the suggestion that the hippocampus was the locus of a cognitive map in the mammalian brain (O'Keefe and Nadel, 1978). The hippocampus is thought to have a broader mnemonic function in higher mammals like monkeys and humans (Scoville and Milner, 1957), such as episodic memory (Kinsbourne and Wood, 1975; O'Keefe and Nadel, 1978; Tulving, 1983). The extensive electrophysiological data available from studies in freely moving rats raise the possibility that the functional properties of the neural machinery of the hippocampus might be understood in the spatial domain, which could also shed light on other putative hippocampal functions for which there are fewer neuronal data.

In this issue of *The Journal of General Physiology*, Muller et al. (1996) present a model in which the synaptic efficacies of the recurrent collaterals in the CA3 region of the hippocampus form a "cognitive graph." They initially demonstrate that CA3 has properties adequate to support a directed connected graph, and they go on to outline their model. In this model, the net resistance of the synaptic connection between two place cells increases monotonically with the time taken to move between the corresponding firing fields ("place fields"), and, hence, on their separation. The graph is built as the animal moves around the environment and long-term potentiation (LTP) causes synapses between place cells with overlapping place fields to strengthen much more than synapses between cells with well separated fields, because of the increased chance of both cells firing within the LTP-permissive time window. Once the graph is formed, such that the connections between

cells within it effectively code for the distance between their place fields, it can be used as follows. If one place cell represents a start location and another represents the goal location, then the chain of place cells between them that involves the strongest synapses (i.e., path of minimum resistance through the network) corresponds to a good route between the start and goal locations (traced out by the place fields along the chain). How this route might be read by the rest of the brain is not explicitly simulated.

This simple model has many attractive features. For example, rats exhibit "latent learning," in which they clearly learn about the layout of an environment in the absence of goals or motivation, helping them to navigate once a goal is provided. This is reflected in the way in which the cognitive graph is automatically constructed during exploration. Similarly, the ability to perform detours and shortcuts exemplifies the type of navigation impaired by hippocampal lesions and can be naturally incorporated in the model. Detours occur because placing an obstacle on or near a place field tends to inhibit the firing of the place cell, thus taking that cell out of the path-planning process. Shortcuts can be found because, while place cells with fields on either side of a barrier will tend to have weak synaptic connection because of the time taken to travel from one field to the other, the rat will tend to explore the area if the barrier is removed and will pass directly from one to the other, causing the connection to strengthen. Another strength of the model is that the directional modulation of place cell firing that occurs when a rat is constrained to run along restricted paths does not destroy the model. However, directional firing would tend to exaggerate the problem of the animal's exploratory behavior, biasing the representation of distance in the graph, for example, so as to underestimate distances along familiar routes compared with unfamiliar ones (see below).

The cognitive graph model appears to work well; indeed, it is reminiscent of a successful path-planning scheme based on "resistive grids" (e.g., see Connelley et al., 1990) that has been extensively studied in robotics

and implemented in neural network form (Tarassenko et al., 1991; Glasius et al., 1995). The essence of this scheme is as follows. Accessible locations in an environment (i.e., points not corresponding to an obstacle) are represented as the nodes in a locally connected grid, and each node is connected to its neighbors on the grid by resistors of equal value. If a current source is placed at the node representing the start location and a sink is placed at the node representing the goal, then the path of maximum current flow (i.e., the path of minimum resistance, as in the cognitive graph) provides a route from start to goal avoiding the boundaries of the accessible environment. The cognitive graph model could be viewed as a way in which this type of scheme could be implemented in CA3 by the action of LTP.

Perhaps the most novel aspect of this model is the postulate that the map itself is stored in the synaptic efficacies. Other models have concentrated on how place cell firing could produce an output to guide behavior that is coded in the firing rates of neurons (Burgess et al., 1994; Brown and Sharp, 1995). The authors' concentration on synapses has advantages in terms of simplicity, and not least for explicitly suggesting a specific role for the CA3 recurrent collaterals (see also Blum and Abbott, 1995) other than as a general associative memory system (e.g., see Marr, 1971), for which direct evidence is more difficult to find. However, if the information in the cognitive graph is to find its way out of the hippocampus and affect behavior, a plausible mechanism for translating the synaptic information into neuronal firing is crucial. In the remainder of this article, we consider the output problem and other issues relevant to the cognitive graph and to hippocampal modeling in general.

Setting Correct Synaptic Resistances

Muller et al. (1996) examine the effects of using different functions relating distance to synaptic resistance in some detail, but they only briefly allude to the difficulty with which such relations could be achieved in the animal. The fundamental problem here is the need to translate information acquired during temporally organized behavior (exploration) into the static structure of the synaptic resistances. One such problem is distortion of the graph by nonuniform exploration. If the graph is constructed by LTP during exploration, then synapses between cells with fields along familiar routes would be strengthened more than those between cells with fields (equally far apart) on an unfamiliar route. Thus, distances in familiar places, or along familiar routes, would be underestimated, biasing behavior towards those places and routes. A second problem is how to deal with nonuniform speeds, i.e., how two runs through the same set of place fields at different speeds translate into the same synaptic changes. The cognitive graph

model appears to duck this problem by assigning different paths through the graph to the same environmental path, such that fast runs involve place cells with large fields and slow runs involve those with small fields. However, it is not clear how this solves the problem, or how different sets of cells could be selected as a function of speed (which will normally vary during each run). In terms of producing constant amounts of synaptic change in the face of speed variations (i.e., different amounts of time spent in each place field), models based on firing rates (in which the rate increases with speed) or the phase of firing with respect to some clock signal (in which phase shifts more quickly with speed) would appear to offer more elegant solutions. In this section we, like the authors, have naturally assumed LTP to be the mechanism behind the construction of the cognitive graph; however, perhaps a little caution is necessary in light of recent evidence that learning the location of the hidden platform in a water maze may not require LTP, at least in some parts of the hippocampus (Nosten-Bertrand et al., 1996) or in rats that have prior experience of the task (see Saucier and Cain, 1995; Bannerman et al., 1995).

The Role of Exploration

The ability of the cognitive graph model to perform shortcuts and detours relies heavily on the acquisition of information during the exploration elicited by the change to the environment. For example, shortcuts become possible only after exploration of the entire extent of a newly opened path or unexplored region, after which the place cells along the new path become possible solutions of the minimum-resistance graph search. The original cognitive map theory (O'Keefe and Nadel, 1978) suggested that this exploration was driven by a mismatch function in the map itself, signaled by the increased firing of some place cells ("misplace cells") that accompanies the myostatial sniffing elicited by the introduction or removal of an object in the place field. This was partially motivated by evidence that hippocampal damage leads to a loss of exploration of novel environments or in response to changed locations of familiar objects in a familiar environment (O'Keefe and Nadel, 1978, chapter 6; Xavier et al., 1990; Thinus-Blanc et al., 1991). This aspect of exploration is not considered in the cognitive graph model, nor, to be fair, in the other models of hippocampal navigation of which we are aware.

Evidence from the early maze learning literature suggested that rats learned the direction to the goal independently of specific paths through the maze (e.g., Dashiell, 1925; Tolman, 1948; O'Keefe and Nadel, 1978, chapter 7), and that this "direction sense" aided the solution of shortcut or detour problems. Thus, given several novel paths from which to choose, the

one leading in the direction of the goal would be most likely to be chosen. We would expect the direction and distance information relating to goals, independent of the paths taken through the maze, to be derivable from the map. For example, the direction and proximity of previously visited goals or obstacles might be signaled by the population vector of the firing rates of sets of output cells that are driven by the place cells (Burgess et al., 1994).

The Nature of Place Fields

The cognitive graph model takes the phenomenology of place cells as given and does not address how or why a given cell fires in a particular location in a given environment. This is understandable, given the focus of the paper, but some of the assumptions that have significant implications for the cognitive graph deserve further consideration. Recent work in our laboratory (O'Keefe and Burgess, 1996) provides evidence that place fields are formed by the thresholded summation of two or more Gaussians peaked at fixed distances from walls in particular allocentric directions. This has several implications relevant to the assumptions of the cognitive graph: (a) that the place fields are dependent on the directional system, and that the absence of a major effect of postsubicular lesions on place fields may be due to a hippocampal input from an earlier stage of the system, perhaps from the thalamus; (b) that the place fields in an environment may not be fixed into a rigid structure. We have seen clear cases in which changing the aspect ratio of rectangular environments causes place fields to move relative to each other (while maintaining fixed distances to different walls). These changes in field separations would require the synaptic resistances of the cognitive graph to be able to increase and decrease accordingly; (c) that the effect on a place field of introducing a barrier might not be simply to suppress its firing, but might be better thought of as equivalent to providing another environmental wall of similar orientation. This would often result in shifting the component of the field that was a fixed distance to the wall behind the barrier far enough away to reduce the cell's firing rate severely, but should occasionally lead to a field being pushed away from the barrier, or to a cell with an unused short-distance input increasing its firing rate. Finally, some of the fields in this recent experiment stretched into more than one separate component as the environment expanded, and cells with double fields have also been found in many experiments. It is not clear which distances would be stored in the resistances of connections with these cells.

Reading the Map

Perhaps the biggest problem with the cognitive graph theory is the absence of a mechanism for reading out

the optimal paths it generates so that they could be used in path planning. One suggestion is that there are "goal" or "destination" cells in CA3, i.e., cells that fire at the start of a run (or possibly throughout it) according to the eventual destination rather than the current location. Together with the activation of place cells at the animal's current location, the firing of these cells could enable activation of the place cells along the optimal path via the strong synaptic connections they make with them, which could in turn enable some unspecified system to guide the rat along the physical path formed by their place fields. However, these destination cells have been sought by one of us for some years now, without success. Speakman and O'Keefe (1990) specifically searched for them in a four-armed maze experiment in which starts, goals, and environmental cues were explicitly manipulated, and they failed to find any. The goal cells reported in the experiments of Gothard et al. (1996) do not fit the bill, as they always fire at a fixed (relatively close) distance to the goal.

A second problem with the read-out of the cognitive graph is that the paths must be identified by the firing of place cells. Since the most obvious correlate of place cell firing is the current location, the authors have suggested that the read-out might occur during the sharp waves seen in LIA (i.e., large irregular activity in the electroencephalogram, as opposed to the regular "theta" oscillation observed during locomotion). A problem here is that sharp waves are prominent during slow wave sleep and consummatory behaviors as well as quiet sitting, behaviors not necessarily found at the start of a run. Synchronous activity that occurs during these bursts is more likely to be correlated with the recent experience of the animal than with its future plans (Skaggs and McNaughton, 1996). Vicarious trial and error (VTE) at a choice point is also not a good candidate for this type of read-out, since in our experiments the only place cells active during VTE are those with fields at the choice point. An alternative mechanism for generating a path from place cell firing is suggested by the finding that some place cells apparently code for locations slightly ahead of the animal (Muller and Kubie, 1989). The location represented by these cells might correspond to the next step along the desired path and so could somehow be compared with the representation of the current location so as to produce a direction of travel. Indeed, it has been suggested that the CA3 recurrent collaterals serve to shift the firing of place cells forward along familiar routes for this reason (Blum and Abbott, 1995).

In summary, a cognitive map of any type must have a read-out mechanism, and several theories posit the existence of destination or goal cells (e.g., O'Keefe, 1991). However, these types of cells are unlikely to be found in the hippocampus proper. This has forced previous sim-

ulations of hippocampal navigation to invoke cells in or downstream of the subiculum (Burgess et al., 1994) or in the nucleus accumbens (Brown and Sharp, 1995). Experimental demonstration of cells with this type of property would represent a major advance in the theory of navigation.

REFERENCES

- Bannerman, D.M., M.A. Good, S.P. Butcher, M. Ramsay, and R.G.M. Morris. 1995. Distinct components of spatial learning revealed by prior training and NMDA receptor blockade. *Nature (Lond.)* 378:182–186.
- Blum, K.I., and L.F. Abbott. 1995. Functional significance of long term potentiation in recurrent networks. *Int. J. Neural Systems* 6: 25–31.
- Brown, M.A., and P.E. Sharp. 1995. Simulation of spatial learning in the Morris water maze by a neural network model of the hippocampal formation and nucleus accumbens. *Hippocampus* 5: 171–188.
- Burgess, N., M. Recce, and J. O'Keefe. 1994. A model of hippocampal function. *Neural Networks* 7:1065–1081.
- Connolly, C.I., J.B. Burns, and R. Weiss. 1990. Path planning using Laplace's equation. *Proc. 1990 IEEE Int. Conference on Robotics and Automation*. 2102–2106.
- Dashiell, J.F. 1930. Direction orientation in maze running by the white rat. *Comp. Psychol. Monogr.* 7:1–72.
- Gladius, R., A. Komoda, and S. Gielen. 1995. Neural-network dynamics for path planning and obstacle avoidance. *Neural Networks* 8:125–133.
- Gothard, K.M., W.E. Skaggs, and B.L. McNaughton. 1996. Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *J. Neurosci.* 16:823–835.
- Kinsbourne, M., and F. Wood. 1975. Short-term memory and the amnesic syndrome. In *Short-Term Memory*. D. Deutsch and J.A. Deutsch, editors. Academic Press, New York. 257–291.
- Marr, D. 1971. Simple memory: a theory for archicortex. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* 262:23–81.
- McNaughton, B.L., C. Barnes, and J. O'Keefe. 1983. The contribution of position, direction, and velocity to single unit activity in the hippocampus of freely moving rats. *Exp. Brain Res.* 52:41–49.
- Muller, R.U., and J.L. Kubie. 1989. The firing of hippocampal place cells predicts the future position of freely moving rats. *J. Neurosci.* 9:4101–4110.
- Muller, R.U., J.L. Kubie, and J.B. Ranck, Jr. 1987. Spatial firing patterns of hippocampal complex spike cells in a fixed environment. *J. Neurosci.* 7:1935–1950.
- Muller, R.U., M. Stead, and J. Pach. 1996. The hippocampus as a cognitive graph. *J. Gen. Physiol.* 107:663–694.
- Morris, R.G.M., P. Garrard, J.N.P. Rawlins, and J. O'Keefe. 1982. Place navigation impaired in rats with hippocampal lesions. *Nature (Lond.)* 297:681–683.
- Nosten-Bertrand, M., M.L. Errington, K.P.S.J. Murphy, Y. Tokugawa, E. Barboni, E. Kozlova, D. Michalovich, R.G.M. Morris, J. Silver, C.L. Stewart, et al. 1996. Normal spatial learning despite regional inhibition of LTP in mice lacking Thy-1. *Nature (Lond.)* 379:826–829.
- O'Keefe, J. 1991. The hippocampal cognitive map and navigational strategies. In *Brain and Space*. J. Paillard, editor. Oxford University Press, Oxford. 273–295.
- O'Keefe, J., and N. Burgess. 1996. Geometric determinants of the place fields of hippocampal neurones. *Nature (Lond.)*. In press.
- O'Keefe, J., and J. Dostrovsky. 1971. The hippocampus as a spatial map: preliminary evidence from unit activity in the freely moving rat. *Brain Res.* 34:171–175.
- O'Keefe, J., and L. Nadel. 1978. *The Hippocampus as a Cognitive Map*. Clarendon Press, Oxford.
- O'Keefe, J., L. Nadel, S. Keightley, and D. Kill. 1975. Fornix lesions selectively abolish place learning in the rat. *Exp. Neurol.* 48:152–166.
- Saucier, D., and D.P. Cain. 1995. Spatial learning without NMDA receptor-dependent long-term potentiation. *Nature (Lond.)* 378: 186–189.
- Scoville, W.B., and B. Milner. 1957. Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg.* 20:11–21.
- Skaggs, W.E., and B.L. McNaughton. 1996. Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science (Wash. DC)* 271:1870–1872.
- Speakman, A., and J. O'Keefe. 1990. Hippocampal complex spike cells do not change their place fields if the goal is moved within a cue controlled environment. *Eur. J. Neurosci.* 7:544–555.
- Sutherland, R.J., B. Kolb, and I. Wishaw. 1982. Spatial mapping: definitive disruption by hippocampal or medial frontal cortical damage in the rat. *Neurosci. Lett.* 31:271–276.
- Tarassenko, L., M. Brownlow, G. Marshall, and J. Tombs. 1991. Real-time autonomous robot navigation using VLSI neural networks. In *Neural Information Processing Systems 3*. R.P. Lippmann, E.J. Moody, and D.S. Touretzky, editors. Morgan Kaufmann, San Mateo, CA. 422–428.
- Thinus-Blanc, C., E. Save, M.-C. Buhot, and B. Poucet. 1991. The hippocampus, exploratory activity, and spatial memory. In *Brain and Space*. J. Paillard, editor. Oxford University Press, Oxford. 334–352.
- Tolman, E.C. 1948. Cognitive maps in rats and men. *Psychol. Rev.* 55:189–208.
- Tulving, E. 1983. *Elements of Episodic Memory*. Clarendon Press, Oxford.
- Xavier, G.F. 1990. Rats with dorsal hippocampal lesions do react to new stimuli but not to spatial changes of known stimuli. *Behav. Neural Biol.* 54:172–183.