

Investigating neural representations: the tale of place cells

William Bechtel

Received: 23 October 2013 / Accepted: 26 April 2014 / Published online: 20 May 2014
© Springer Science+Business Media Dordrecht 2014

Abstract While neuroscientists often characterize brain activity as representational, many philosophers have construed these accounts as just theorists' glosses on the mechanism. Moreover, philosophical discussions commonly focus on finished accounts of explanation, not research in progress. I adopt a different perspective, considering how characterizations of neural activity as representational contributes to the development of mechanistic accounts, guiding the investigations neuroscientists pursue as they work from an initial proposal to a more detailed understanding of a mechanism. I develop one illustrative example involving research on the information-processing mechanisms mammals employ in navigating their environments. This research was galvanized by the discovery in the 1970s of place cells in the hippocampus. This discovery prompted research in what the activity of these cells represents and how place representations figure in navigation. It also led to the discovery of a host of other types of neurons—grid cells, head-direction cells, boundary cells—that carry other types of spatial information and interact with place cells in the mechanism underlying spatial navigation. As I will try to make clear, the research is explicitly devoted to identifying representations and determining how they are constructed and used in an information processing mechanism. Construals of neural activity as representations are not mere glosses but are characterizations to which neuroscientists are committed in the development of their explanatory accounts.

Keywords Mechanistic explanations · Representations · Spatial navigation · Place cells · Grid cells

W. Bechtel (✉)

Department of Philosophy, Center for Chronobiology, and Interdisciplinary Program in Cognitive Science, University of California, La Jolla, 0119, San Diego, CA 92093-0119, USA
e-mail: bechtel@ucsd.edu

1 Introduction

The concept of representation figures centrally in philosophical discussions of neuroscience. This is appropriate since neuroscientists often employ representational vocabulary to characterize various neural processes (the rate or pattern of action potentials, synchronized electrical potentials, etc.). A strategy neuroscientists have employed with great success in attempting to understand the mechanisms that underlie cognitive abilities is to identify cells in which the rate of action potentials increases in response to specific stimulus conditions. They then construe such neurons as representing those features in the environment whose presence is correlated with the increased firing and attempt to understand how subsequent neural processing utilizes representations that stand in for those features of the environment in guiding behavior. Questions that philosophers (see, for example, [Haselager et al. 2003](#)) are prone to ask are: Is the characterization of neural activity as representational critical to the research the scientists are pursuing? Or is the construal of neural activity as representational only an interpretation imposed and used by scientists?

Philosophers of the cognitive and neurosciences have advanced a variety of arguments against accounts of neural mechanisms that characterize them as employing representations. I will focus on arguments, especially those advanced by Ramsey and Egan, that accept that neural processes might be interpreted as representations but that this only provides a way for researchers to identify these processes or relate them to the tasks they perform. My strategy will be to show, through a detailed case study of research on rodent navigation, that the research these scientists have pursued was directed at determining which neural processes serve as representations and what their contents are. Independent of that objective, there would be no reason to conduct those types of research.

Other critics of identifying representations in the brain challenge the whole project of understanding the brain as an information-processing mechanism that employs representations. Critics such as [Gelder and Port \(1995\)](#) and [Chemero \(2000, 2009\)](#) offer an alternative explanatory framework that is not mechanistic in character. Inspired by dynamical systems theorists such as [Kelso \(1995\)](#), [Jirsa et al. \(1998\)](#), they advocate the development of dynamical models that do not employ mechanistic decomposition to identify components that carry information but rather seek mathematical relations between variables characterizing the brain, body, and environment that account for the observed data. This is not the place for a systematic response to these arguments.¹ One should note, though, that not all dynamical accounts reject mechanistic decomposition. Many researchers turn to dynamical models to understand how mechanisms, whose parts have been identified as representing information, will behave (for examples of such uses of dynamical models in the case of rodent navigation, see [Burgess and O’Keefe 2011](#); [Navratilova et al. 2012](#)). The relevant alternatives to explanations invoking representations are dynamical systems models in which the components described in variables are identified without any appeal to their representational function. The conflict here is between radically different research programs and such

¹ I have address van Gelder’s arguments in [Bechtel \(2008\)](#) and Chemero’s arguments in [Kaplan and Bechtel \(2011\)](#).

conflicts, as [Kuhn \(1970\)](#), [Lakatos \(1970\)](#), and others have argued, are typically settled by the ability of each group of practitioners to make what they perceive as progress. To date those viewing the brain in information processing terms still see themselves as making progress in understanding neural mechanisms as processing information. Like all research programs, theirs may eventually stall and be replaced by an alternative that rejects understanding brain activity in representational terms. My analysis is not directed at predicting future outcomes, but rather articulating the role representations play in this currently thriving research program.

2 Are appeals to representations mere glosses?

Before turning to the case study and what it shows about how questions about representation have directed the research on neural mechanisms involved in rodent navigation, I consider the arguments advanced by Ramsey and Egan that these neural processes are not true representations but only construed as such in the researchers' interpretations of their research. Given the prominence of Ramsey's arguments against representations, I will first situate the discussion in the context of the responses that have been raised to his argument. I then focus on Egan's contention that appeals to representations are mere glosses, against which I will argue throughout by showing how research on rodent navigation has been directed at determining what are the representations and what they represent.

[Ramsey \(2007\)](#) develops a sustained argument that what are treated as representations by neuroscientists and others (especially those advancing neural network models) are not true representations. On his account, representations in the proper sense of the term were invoked in classical symbolic models in cognitive science. These accounts posited what he terms, following [Swoyer \(1991\)](#), *structural representations*. Such representations preserve the structure of the represented domain for which they stand in and play a causal role in the ability of the cognitive system to generate behaviors (see also Cummins's, 1989, notion of simulation representation). Ramsey contends that what are being called *representations* in more recent accounts of cognition are only *receptor* or *tacit* representations—responses in the brain that researchers might label in terms of the stimuli that cause them. Their subsequent role in the system can be characterized purely in terms of mechanical processing (e.g., as a sequence of switches each of which activates the next) without considering their representational function. Researchers trying to understand neural network models might use techniques such as cluster analysis to show how information is encoded in the activity of units or the weights in the network, but on Ramsey's account this serves simply to aid researchers in tracking the processing of the network. The network does not use the structural correspondence to what is represented in its processing.

Several philosophers have challenged Ramsey's deployment of the distinction between structural and receptor representations. Some accept the distinction but have argued that neural network models in fact employ structural representations. Garzón and Rodríguez (2009, p. 308), for example, argue that cluster analysis, when used to analyze connectionist networks, reveals "isomorphic relations between the components of the modeling nets and the constituents of the target domain" and that it is

the architectural relations between these components that explain the performance of these networks. Shagrir (2012, p. 527) focuses on recurrent networks, which he contends are far more characteristic of current neural network modeling than feedforward networks that were the focus of Ramsey's analysis, and argues they are used "with the explicit aim of positing the notion of an internal model." He supports this claim by analyzing research modeling the human oculomotor memory using recurrent networks that implement continuous attractors. He shows both how location is represented in the network and how these representations are updated by performing the computational operation of integration on the representations.²

In contrast, Morgan (2014) challenges the use of the distinction itself, arguing that what might seem to be receptor representations also meet the criterion for structural representations. In developing this argument, Morgan focuses not on computational models, but on the neural systems being modeled. He develops examples of neural systems that seem to involve receptor representations, such as thermosensors in fruit flies, but argues that these sensors provide structural representations by showing that serve to guide appropriate behavior with respect to what is represented. Moreover, it is not just the scientist who is interpreting these sensors as representing temperature but the parts of the organism that use them to control behavior.³

Like Morgan, I will focus on research that is directed at discovering how the brain represents information in ways that can guide behavior, not the efforts to model the brain.⁴ I will turn, though, to another aspect of research directed at representations in the brain that has not been sufficiently discussed—how research directed at representational vehicles and their contents develops over time. I will show that treating brain processes as representations is foundational to these research projects as they have been pursued. The research pursuits are focused on identifying what neural processes serve as representational vehicles and especially determining what they represent. The research efforts are designed to answer these questions and it would be difficult to understand why the researchers pursue these projects if their goal was not to identify representations and determine what they represent.

In examining how identifying representations figures in the development of neuroscience research, my main focus will be on Egan's (2010) contention that treating brain processes as representations and understanding them in terms of their content is simply a gloss employed by scientists. Egan defends the view that she attributes to

² An important part of Shagrir's argument is that in determining what computation the oculomotor system is performing the researchers turned to the relation between the position and velocity of the eyes that the system was representing—it is because the relation between the represented properties is integration that the brain needs to perform integration over representations.

³ In making this case Morgan also draws attention to a condition in Dretske's (1988) account of representations, that the representation be used in the "production of movements whose successful outcome depends on what is being indicated" (p. 70). Below I contend that representations arise in systems that have the task of controlling other systems. This makes their use in the control of behavior central to accounts of representations.

⁴ Morgan shows that some of the objections Ramsey has raised about representations in current neuroscience are more properly concerned with mental representations, where such representations figure in the mental models of agents. For neuroscience, the relevant notion of representation needs to be dissociated from the role representations play in mental models and treated as processes within a processing system that stands in for their contents without the organism being conscious of them.

Chomsky (1995) that content-based interpretations of cognitive processes serve only as an “informal presentation” of cognitive theories. Drawing on Marr’s (1982) characterization of three levels of analysis in vision, she argues that the computational level treats the cognitive system as performing a mathematically characterized computation; appealing to content-based glosses such as filtering the visual image only shows how the theory is able to explain cognitive abilities. Shagrir (2010; see also Shagrir and Bechtel 2011) challenge Egan’s interpretation of Marr, arguing that a central task of Marr’s computational level was to explain why performing the computation is appropriate given the task posed by the world. The ability of internal states to represent relevant features of the world plays an important role in explaining the appropriateness of specific computations. My contention in this paper is related but distinct: I am arguing that much neuroscience research is in fact directed at determining which neural processes are content bearers and understanding how they represent what they do. Content characterizations are not glosses on the research; the goal of the research is to determine what content the representations have.

Philosophical discussions of representations often distinguish representational vehicles and the contents that they representation. Although they are conceptually distinct, neuroscience research often integrates the two—using investigations of content to help identify the vehicles. Thus, an early and integral step in the investigation of how specific information is processed within organisms appeals to representational content to determine representational vehicles. Initial characterizations of the vehicles and attributions of content are then both subject to revision as more vehicles are discovered and the processing mechanisms that generate the relevant activity and respond to it are identified. What is especially important is that such additional inquiry is inspired and guided by the initial attributions of representational content and directed at fleshing out the account. The attribution of content is a first step in articulating an account of a mechanism for processing information. Without this initial assignment of representational content, researchers would not be able to formulate the hypotheses that guide subsequent research. For example, only once a given population of neurons is hypothesized to represent information in a particular manner can researchers formulate hypotheses about possible sources and uses of that information, including hypotheses about what other representations must exist and the processes through which these representations interact in generating behavior. Identifying representations thus undergirds the inquiry—it is not merely a gloss attributed at the end.⁵

My contention, then, is that identifying representations is an important aspect of neuroscientists’ quest to explain mental phenomena by identifying and characterizing the information-processing mechanisms responsible for them. The quest to explain phenomena by identifying mechanisms is widespread in many fields of biology, not just neuroscience, and an important step in developing explanations of mechanisms is

⁵ The language I use in this paragraph to describe the project of identifying representations is very similar to that which McCauley and I (Bechtel and McCauley 1999; McCauley and Bechtel 2001) used in discussing identity claims in science—they are proposed at the outset of inquiry and by the time a research endeavor has developed around them, revising and elaborating on them, it would seem perverse to the scientists to propose that the relation in question was mere correlation, not identity. In fact claims about representations can be viewed as instances of identity claims since what researches are doing is identifying constituents of mechanisms as representations.

to decompose them into parts and operations that, when appropriately organized and orchestrated, enable the mechanisms to produce the phenomena in question (Bechtel and Richardson 1993/2010; Bechtel and Abrahamsen 2005; Machamer et al. 2000). In many fields of biology, the phenomena under investigation involve the generation, degradation, or transformation of some identifiable entity—fermentation converts sugar into alcohol while liberating energy that is captured in ATP. Some of the phenomena for which parts of the brain are the relevant mechanism are different—they involve regulating or controlling other organs within the organism that enable the organism to coordinate its behavior with distal features of the environment.⁶ The brain is functioning as a control system and control processes require information about the plant and environment (including on occasion what conditions were like in the recent to distant past) that the controller can use in developing plans for action (that may not be executed until some time in the future). An important part of what neuroscientists want to understand about the brain is how it contributes to controlling processes within the organism or its behavior in the external world. It is in this context that identifying representations and their content becomes critical—it is as they represent entities and processes external to the brain that mechanisms within the brain figure in these control processes (Bechtel 2011).⁷

In this paper I will employ research on the information-processing mechanisms mammals employ in navigating their environments as an illustrative example. This research was galvanized by the discovery of neurons (place cells) that generate action potentials primarily when the organism is in a particular region of its local environment. Action potentials of these neurons were interpreted as representing that location. The identification of these neurons raised questions about how place representations are constructed in the relevant population of neurons and how they contribute to navigation, questions researchers tried to address by manipulating factors that altered the behavior of these neurons. The research also led to the discovery of a host of other types of neurons—grid cells, head-direction cells, boundary cells—that encode other spatial information that is used in performing navigational tasks. Although this research is still ongoing, and one can expect many more discoveries and revisions of the current understanding in the future, it is sufficiently developed that it has identified many of the parts and operations involved in the mechanism of rodent navigation. By examining it we can recognize how identifying representations is foundational to such neuroscience research. As I will try to make clear, the research is explic-

⁶ This is not to ignore neural phenomena that involve local transformations within the brain that are also of considerable interest to neuroscientists. In addition to ongoing metabolic activities and processes of gene expression, there are operations involving electrical potentials resulting from movements of ions across membranes which can sometimes give rise to action potentials in the absence of external stimuli. Some of this endogenous activity, such as synchronized oscillations in electrical potentials, may be extremely important to the brain's capacity to execute cognitive operations when it receives sensory input (Abrahamsen and Bechtel 2011).

⁷ Processing information so as to coordinate responses to conditions inside or external to the organism is not unique to organisms with brains. Bacteria perform a vast array of information processing functions using chemical rather than neural signaling, and this has led a number of researchers to refer to bacterial cognition (see, for example, Shapiro 2007). Plants also use chemical signaling to execute a broad range of behaviors that has inspired a field some refer to as plant neurobiology (for discussion, see Garzon and Keijzer 2011).

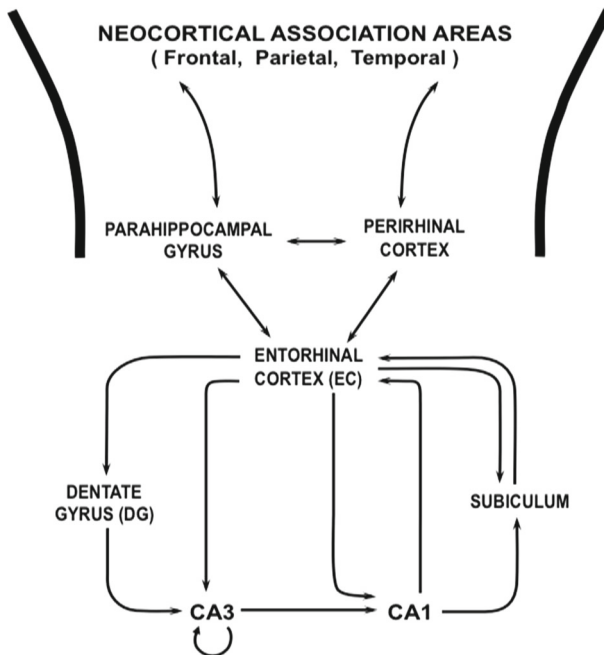


Fig. 1 Schematic representation of the organization of the hippocampus and related structures in the medial temporal lobe

itly devoted to identifying representations and determining how they are constructed and used in information-processing mechanisms that control behavior. Characterizing neural processes as representations is not viewed as just a convenient way of talking about brain processes. The research is predicated on these processes being representations; the explanatory tasks they devote themselves to are identifying those neural processes that are representations, figuring out what their content is, and how these representations are then used in controlling behavior.

3 The discovery of place cells

The hippocampus has long been a target region for neuroscience research, due in large part to its distinctive architecture—comprising regions with distinct types of neurons and patterns of connectivity.⁸ The hippocampus lies deep within the medial temporal cortex of the mammalian brain. As indicated in Fig. 1, it receives inputs from and sends its outputs to the entorhinal cortex;⁹ the entorhinal cortex in turn is connected reciprocally

⁸ See Craver (2003) for a discussion of the research on the hippocampus that led to the discovery of long-term potentiation as a laboratory technique before it, and the hippocampus more generally, became associated with learning and memory.

⁹ Miller and Best (1980) established the necessity of input from entorhinal cortex for hippocampal function by showing that lesions to the entorhinal cortex impaired the navigational abilities of rats and the responsiveness of hippocampal place cells.

cally with the perirhinal cortex and perahippocampal gyrus, through which input from and outputs to the neocortex pass. The central components of the hippocampus are connected pairwise into a sequential pathway for information flow: dentate gyrus (DG) → CA3 → CA1 → subiculum. The DG receives input from layer II of the medial entorhinal cortex (MEC) and typically only a few of its neurons fire on a given occasion, producing what is known as sparse firing that is thought to provide a sparse coding of the MEC input. Via mossy fibers, individual DG neurons have a large number of synapses on specific CA3 neurons that enable a single DG neuron to induce an action potential in a target CA3 neuron. CA3 neurons have extensive interconnections between themselves leading researchers to hypothesize that the region functions as an auto-associative memory system that can complete patterns from partial information (Marr 1971). There are more neurons in CA1 than in CA3, and each CA3 neuron projects to a large number of CA1 neurons, suggesting a further processing of the information encoded in CA3. In addition to this indirect pathway, there are also projections directly from MEC (layer III) to CA1; this direct pathway will become relevant later. CA1 can send outputs directly back to MEC (layers V and VI), or route these first through the subiculum. (Although the projections to and back from the hippocampus involve different layers of MEC, the loop is actually closed since there are projections from the deep layers (V–VI) of MEC to the superficial layers (I–III), but not in the opposite direction.)

A wide variety of hypotheses have been advanced as to function of the hippocampus. The appearance of complete anterograde and significant retrograde amnesia in the patient known as HM after resection of his hippocampus in an attempt to reduce the effects of epilepsy (Scoville and Milner 1957) led human researchers to focus on its function in the development of long-term declarative memories. Animal research, especially in rodents, instead, focused on spatial navigation. By 1970 it was known that rats with hippocampal damage exhibit a variety of deficits, including in spatial navigation, prompting O’Keefe and Dostrovsky (1971) to implant electrodes into regions DG, CA1 and CA4 of the hippocampus to determine the roles of these regions in spatial navigation. They found that 8 of the 76 neurons from which they were able to record “responded solely or maximally when the rat was situated in a particular part of the testing platform facing in a particular direction” (p. 172). This initial “Short Communication” differs in interesting ways from subsequent reports of place cells: O’Keefe and Dostrovsky report that firing of these neurons occurred only when the rat was oriented in a specific direction (as shown in Fig. 2) and “was simultaneously lightly but firmly restrained by a hand placed over its back with thumb and index finger on its shoulder and upper arm” (p. 172). Despite the need for specific orientation and tactile stimulation to elicit activity from these neurons, O’Keefe and Dostrovsky reached a bold conclusion that has inspired several decades of subsequent research:

These findings suggest that the hippocampus provides the rest of the brain with a spatial reference map. The activity of cells in such a map would specify the direction in which the rat was pointing relative to environmental land marks and the occurrence of particular tactile, visual, etc., stimuli whilst facing in that orientation. The internal wiring of the hippocampus, on this model, would be such that activation of those cells specifying a particular orientation together with a signal indicating movement or intention to move in space (hippocampal

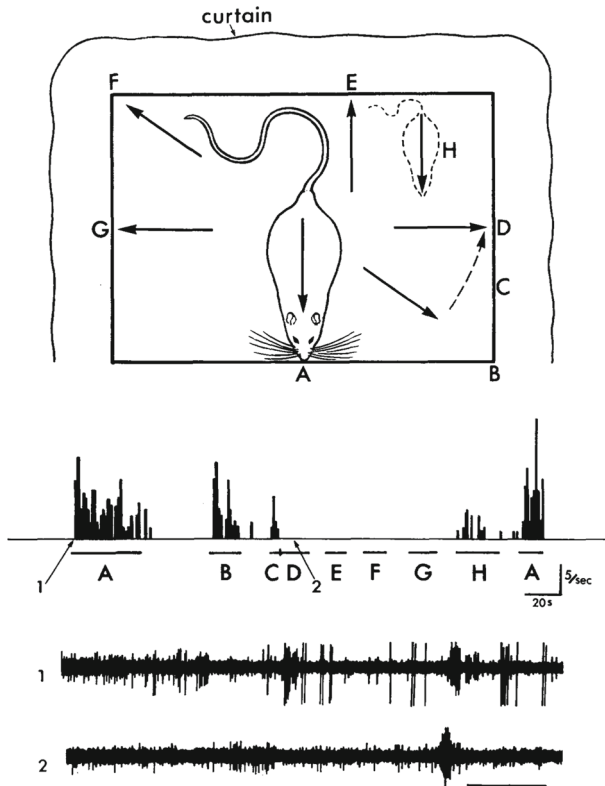


Fig. 2 The results of recording from a single neuron in CA1 of the hippocampus as the rat moved between locations in the enclosure marked with *letters*. The histograms in the *middle* indicate the spikes recorded while the animal was restrained in the location (*letters* represent time at the location and the *lines* above them are the periods when the rat was restrained at the location). The *bottom two lines* indicate the spikes recorded while in location A and D respectively. Reprinted with permission from Elsevier from O’Keefe and Dostrovsky (1971)

Θ and Θ -related movement units) would tend to activate cells specifying adjacent or subsequent spatial orientations. In this way, the map would ‘anticipate’ the sensory stimuli consequent to a particular movement (p. 174).

As noted above, only 8 of the 76 from which O’Keefe and Dostrovsky recorded responded to locations. Of the other cells, they characterized 14 as “arousal” or “attentional” units and 21 as “movement” units based on the stimulus conditions that would elicit responses in them. The latter are the Θ units referred to in the above passage. Shortly thereafter Ranck (1973), showed that these cells produced a regular spike train and increased their firing rate in the presence of theta rhythms (regular electrical oscillations of 6–10 Hz detected with EEG during either voluntary activity or rapid eye movement sleep).

In their initial short communication, O’Keefe and Dostrovsky had assigned no name to the cells that responded to locations. Ranck characterized them as “complex spike cells” because they produced, at least on some occasions, a sequence of 2–7 spikes

with varying amplitude and an interspike interval of 1.5–6 ms. Ranck's objective was to identify behavioral correlates of cells (a project he termed *microphrenology*¹⁰) and one class of complex spike cells he distinguished as approach-consummate-mismatch cells (an example would be a cell that generated complex spikes when the animal approached or drank water, or when it approached but found no water at a normal location). Ranck noted that he had completed his experimental work before becoming aware of O'Keefe and Dostrovsky's report, but concluded these were the same type of cell from which O'Keefe and Dostrovsky had recorded and comments "Perhaps spatial characteristics are the entire basis of firing in these cells. The evidence at present does not allow us to decide" (p. 498).

In a paper reporting a follow-up study, O'Keefe (1976) referred to these cells as *place units* and near the end as *place cells*, the term which became standard. (In this paper he also referred to Θ units as *displace units*.) This paper also reflected important changes both in the methods employed and in the data reported. Methodologically, O'Keefe changed from a strategy of manually situating the rat in an enclosed arena to allowing it to run in an elevated three-arm maze whose sides were open to the surrounding laboratory. He also recorded only from neurons in the CA1 region. In this study he found that 26 cells (out of the 50 from which he recorded) responded primarily to location, with 20 responding when the rat occupied or ran past the appropriate location (which he designated the *place field*). The other 6 responded most strongly when the rat did not find the expected food or water and began exploratory sniffing at the location. Gone from the reports of the data is any reference to the direction the rat was facing (although this was highly restricted by the structure of the maze) or the need to restrain the rat. The only variable that was correlated with firing rate was location and so this was what the researchers treated the cells as representing.

Discovering that the firing of a particular class of neurons depended on the rat occupying a particular place field was sufficient for O'Keefe to characterize them as providing a spatial map that represented where the rat was in the world. The idea that rodents rely on a map in solving navigation tasks was advanced by Tolman (1948) on the basis of behavioral studies showing that rats would follow routes in mazes that led them more directly to their goal than those on which they had been trained. This suggested that rats must have a representation of the spatial layout of their environment,

¹⁰ Ranck provides a reflective discussion of the prospects and challenges of such a project noting several reasons the project might fail: the firing of a neuron "may signal something not directly related to overt behavior, such as drive state, or some idea the rat has, or the blood level of some substance. The firing of the neuron might be part of some internal timing mechanism, or a mechanism in memory retrieval. The firing of the neuron might be significant only in some neural net, and therefore, firing of a single neuron may not be interpretable." With the hippocampus, however, he proposes that the strategy does work, but contends that correlation is not enough—one must determine how the information is transformed: "to be able to apply this approach to hippocampal formation we must know behavioral correlates of almost all inputs and outputs of the system and see what transformations occur." He distinguishes two strategies, one in which researchers seek out conditions in which the neuron fires (during this stage "*the behavior of the neuron shapes the behavior of the experimenter*"), and a second in which a more systematic protocol is employed that also considers its firing frequency and patterns. Ranck's discussion is particularly interesting in light of Ramsey's claims, discussed above, that neuroscience only uses a receptor notion of representation. On Ranck's view, understanding the behavior of neurons requires going beyond identifying it as a receptor and determining how what it represents is affected by the context within the organism.

which he termed a *cognitive map*. Tolman had little to say either about how a mechanism using such a map would work or where it was located in the brain. Drawing on his findings, O’Keefe together with Nadel, wrote the influential book *The hippocampus as a cognitive map* (O’Keefe and Nadel 1978) that, as the title suggests, presented the hippocampus as instantiating Tolman’s hypothesized cognitive map.¹¹ O’Keefe and Nadel viewed the place cells of the hippocampus as providing an allocentric representation of space that, like representations in cartographic maps with which we are familiar, enable animals to determine novel routes. Such navigation with allocentric representations is contrasted with navigation in terms of egocentric directions using landmarks (e.g., walk 100 meters until there is a café on your left, then turn right). This latter ability, they argued, did not depend on the hippocampus). Although place cell representations were viewed as like cartographic maps in being allocentric, they were not viewed as laid out in the same manner. Whereas in a cartographic map the spatial locations between representations correspond, albeit only approximately and with distortions, to the spatial relations between the places represented, this is not true of the map realized in place cells. Rather, functional relations between neurons are presumed to encode the relations between the places represented and to enable the animal to calculate routes between the places represented.

4 Figuring out how place cells represent locations

In advancing the spatial map hypothesis, O’Keefe was clearly construing place cells as representing locations in space. The evidence that action potentials in place cells represent a rat’s location in space—the sounds played on a loudspeaker in response to the spikes from a specific neuron when a rat is in a particular region of its enclosure (or indications of places where the cell has fired onto tracings of the route the rat has followed as in Fig. 3)—seemed compelling to many. But once he construed place cells as representations of locations in space, O’Keefe wanted to determine whether it was place per se or some surrogate that the animal was representing and what enabled it to do so. Together with Conway (O’Keefe and Conway 1978), he posed a set of related questions:

Is [representing a place field] due to something the rat does in the place field or to some environmental factor? If the latter, is the cell responding to a stimulus, or is it signalling more abstract information such as the place itself, as we have previously suggested? How does the cell identify the place? Does it do so on the basis of a special set of cues or will any cue do? (p. 574).

As this passage suggests, the ensuing research project focused not on the proximal mechanism (determining from which cells place cells received their input) but on the

¹¹ In the wake of the initial research recording from hippocampal cells, O’Keefe collaborated with Nadel on a lesion experiment in which they lesioned the major input and output pathway from the hippocampus, the fornix. They found that these rats were unable to learn to locate water that was always at the same location, but showed somewhat improved performance in locating water that was always marked by the same cue (light). This indicated that the rats required the representation of its current location in allocentric space provided by CA1 cells in order to use place information in navigating their environment.




Fig. 3 Locations that elicit action potentials in a place cell (*red dots*) as a rat navigates the path indicated by the *black line*. Reprinted with permission of Annual Reviews from Moser et al. (2008). (Color figure online)

distal stimuli that the action potentials were viewed as representing. The goal was to determine which stimuli enabled the rat to represent a specific feature specified by place cells—the rat's location.

Already in his first paper with Dostrovsky, O'Keefe had recognized that for a given cell to respond to location the animal needed sensory input. Radical changes in sensory input would change the response of a cell, but that response did not seem to be linked to any single feature of the stimulus:

the spatial orientation of two units was eventually disrupted after several radical changes in the environment (such as removing the curtain) but then the rats began to behave (constant exploration) as though they were in a totally new environment. We suspect, but have not proved, that these cells derive their orientation preferences from several equipotential cues, removal of any one of which is insufficient to disrupt the response.

In his 1976 paper O’Keefe determined that place cells continued to respond in darkness, ruling out ongoing visual stimuli as necessary. He also ruled out the necessary reliance on olfactory and tactile cues by replacing a given arm with a substitute of the same dimensions. Although he could not eliminate the possibility that the cell was responsive to a simple sensory feature, he was unable to find one. He explored whether complex groupings of stimuli or the pairing of a stimulus with a behavior could explain the response of these cells and concluded that they could not. Rather, he favored the hypothesis that it was the location in space that mattered and that “input from the navigational system gates the environmental input, allowing only those stimuli occurring when the animal is in a particular place to excite a particular place cell” (p. 107).

In the attempt to address the questions above, O’Keefe and Conway investigated the effects of additional manipulations of the environment. They found that if they shrouded the enclosure in a curtain and then rotated the enclosure and curtain with respect to the external environment, the place cell responses were unaffected, indicating that the environment external to the enclosure was not essential to the activity of place cells. They found, though, that in this arrangement place cell activity was sensitive to changes in the four cues they had mounted on the wall of the enclosure. When all four cues were removed, cells responded to all locations equally and exhibited an increased firing rate. Although no single cue seemed to control the response, two cues often sufficed to elicit normal place cell activity. They concluded: “the place fields can be determined by cues such as lights, sounds, and feels, and are not necessarily dependent on distal cues fixed to the earth’s axis such as geomagnetism” (p. 589).

Given that place cell behavior seems to be influenced by local cues, a pertinent question is what happens when the animal is shifted to a very different environment. O’Keefe and Conway investigated whether the same cells designate places in different environments and, if so, whether the places that elicit responses from the cells have the same spatial relation in different environments. They found that 15 of 34 cells from which they recorded responded to locations on both a platform and a T-maze but could find “no obvious topographic relationship between the place fields in the two environments” (p. 587). Each environment the rat encountered generated a different map in the cell population in the hippocampus.

In referring to place cells as providing a spatial map, O’Keefe and his collaborators clearly assumed that the animal used place cell activity to regulate behavior, but the primary evidence for this was indirect—rats with hippocampal lesions experienced navigational difficulties. Additional evidence, albeit still indirect, was the demonstration that these cells retained their ability to spike in response to locations after the cues from the environment were removed and the rat was forced to delay navigating the maze for a food reward. Not only did [O’Keefe and Speakman \(1987\)](#) find that place cells continued to emit action potentials in such a working memory task, but that the appropriate place cells continued to respond when the animal was forced to take a detour after its release. This is what one would expect if place cell activity is not just a conditioned response to specific stimuli, but encodes the local environment in a spatial map. They conclude: “hippocampal place cells are either the site of the neural changes subserving one form of spatial memory or are ‘downstream’ from that site” (p. 22).

5 Using place cell “Remapping” to study the representation relation

The research reviewed so far sufficed to show that place cells emit action potentials when rats were in specific locations in their enclosure and that, while sensory information was necessary for the rat to locate itself in space, no single cue seemed to determine the response of place cells. Instead, place cells seemed to be primarily carrying information about the animal’s location in its local space. Yet, they must acquire that information in some manner from various sensory cues and the challenge was to figure out how they do so. The basic strategy was to vary the sensory cues available to the animal and determine the effects on the activities of specific place cells. Researchers pursued this objective by altering an existing enclosure in a variety of ways to identify what specific alterations would lead to different responses from place cells.

Muller and Kubie (1987) adopted such an approach with a goal of identifying “a transformation rule for each environmental manipulation, such that the new spatial firing pattern can be predicted from the pattern in the original situation” by systematically varying features of a relatively simple environment. They began with a gray cylindrical chamber 76 cm in diameter with 51 cm walls, with a white cardboard sheet covering 100° of the cylinder’s arc. When they rotated the location of the sheet 90°, the place fields (in all but one case) rotated 90°. However, when they totally removed the cue card, the place fields rotated in an unpredictable manner. When a larger cylindrical chamber was substituted, half the cells had the same relative place fields in both environments, and for most of these the size of the place field was expanded in the larger enclosure. The same results were obtained with two rectangular enclosures of different sizes, but not when a rectangular enclosure was substituted for a cylindrical one. Finally, Muller and Kubie explored the effects of inserting barriers within the enclosure and found that this changed the responses only of those cells in whose receptive fields the barriers were placed (and did so even if the barrier was transparent).

Muller and Kubie introduced the term “remapping” in presenting the results of the barrier experiment in which cells that were previously only weakly active or inactive were recruited to represent part of the space affected by the barrier. The term *remapping* was subsequently generalized to describe how place cells change their response as cues change. By exploring the effects of changing different cues, the research community generated an initially puzzling set of results. Some studies suggested that turning off lights resulted in very little remapping but introducing the rat into a familiar but darkened enclosure produced more general remapping in which half the place cells either stopped responding or responded to a very different place field (Quirk et al. 1990). Although changing the size of the cue card had produced no effects in Muller and Kubie’s initial study, in a subsequent one they found that changing the color of the cue card from white to black led to altered responsiveness of approximately half of the place cells from which they recorded. Moreover, the response was gradual: it took three sessions of 32 min for the cells that changed in response to the change in cue card to exhibit a new regular response pattern (Bostock et al. 1991). In addition, moving a rat from a white to a geometrically identical grey cylindrical enclosure led to widespread remap-

ping. Finally, changing task conditions (Markus et al. 1995) and introducing fear conditioning (Moita et al. 2004) both resulted in changes in place fields or firing rates.

As this brief sketch of the multitude of remapping studies researchers conducted suggests, the results at first did not seem to reveal a clear pattern. Nonetheless, researchers persevered in these investigations since their goal was to determine what place cell activity represented and such research was the way to fulfill that objective. Leutgeb et al. (2005a, b) advanced a conceptual framework that brought some order to these remapping results by distinguishing two properties of place cell activity that served different representational functions. They differentiated rate remapping, in which only the firing rate of place cells changes, from global remapping, in which the place fields that elicit activity also change, and proposed that firing rate and the identity of the cell that fired carried different information: the particular cell that fires codes for the location of the place field, whereas the rate of firing encodes non-spatial information associated with the place. In their study, they found that changing the recording enclosure to one of a different shape or color but in the same room often led to an order of magnitude change in firing rate of the cells but did not alter the place fields to which they responded, whereas identical enclosures in different rooms resulted in changes to both the place field and the firing rate.¹²

Leutgeb et al.'s analysis did not address the process through which global remapping is generated. As researchers approached this question, variations in experimental procedure led to yet more initially confusing sets of findings. As a first step towards characterizing the process, Lever et al. (2002) followed the same place cells, when possible, or different place cells from the same population of cells, as rats were exposed to otherwise identical circular and square enclosures over multiple days. On initial exposure to both boxes, place cells responded to place fields in the same relative positions in the two boxes. The fields to which the cells responded then diverged over successive days. In one rat this happened in just five days whereas in others it took longer. The researchers identified three different patterns of change: the initial emergence of a second place field to which the cell gradually increased its rate of responding, the gradual movement of the place field when in one enclosure, and the gradual diminishing of response in one enclosure. In a subsequent study, Wills et al. (2005) elaborated the strategy for monitoring change in place cell behavior over time. The enclosures they used permitted morphing the shape from square to circular. After familiarizing the rats to both enclosures (resulting in the establishment of place cells with distinctive place fields within each), they morphed the square into a circle through four intermediate forms. They found that in a given rat most place cells switched abruptly at the same

¹² Leutgeb et al.'s findings are, on the face, inconsistent with those of Muller and Kubie noted above in which change from a circular to a rectangular enclosure resulted in global remapping. This may be explained by the fact, that in Muller and Kubie's study the rats were first trained on both enclosures when they were next to each other in a common room and so had presumably developed place codes for each (Colgin et al. 2008). The substitution of one enclosure for another thus elicited the distinct encodings that had already been acquired for each.

intermediate step from responding to the place field they had in the square enclosure to the one they had in the round enclosure.¹³

Employing the same morphing technique but a slightly different procedure, [Leutgeb et al. \(2005a, b\)](#) generated very different results—in their study individual neurons remapped in a stepwise manner that fit a linear or quadratic function.¹⁴ Leutgeb et al. reported significant hysteresis in their primary study, but showed that the gradual transition was equally exhibited when the intermediate enclosures were experienced in random order. One difference between the procedures in the two studies is that the circular enclosure in which [Wills et al.](#) initially trained their rats was of a different color and made of different materials than the square enclosure, whereas Leutgeb et al. began with the morphed version of the square enclosure. The result, as Leutgeb et al. acknowledge, may have resulted in global remapping in [Wills et al.](#)'s study and only rate remapping in their own (see discussion in [Colgin et al. 2008](#)).

A further strategy researchers pursued for determining how stimuli generate place cell activity is to take advantage of a phenomenon [Muller](#) labeled *partial remapping* ([Muller et al. 1991](#)). In partial remapping, different sets of place cells remap in response to specific changes in different cues. For example, reorienting the local enclosure in its larger environment results in some but not all place cells remapping ([Zinyuk et al. 2000](#)). [Skaggs and McNaughton \(1998\)](#) explored a situation in which two identical enclosures were connected by a passageway, and found that some place cells behaved the same in both enclosures while others behaved differently depending on which enclosure the rat was in.

The study of remapping has provided one of the main avenues for studying place cell representations in the hippocampus. The systematic changes in both the place fields and firing rates of place cells to in response to changes in stimuli were pursued by the investigators in their attempt to determine how different properties of the vehicles, cell identity and its firing rate, encoded different information. The motivation for performing these and many other studies was to pin down exactly what changes in stimuli result in specific forms of place cell remapping. It is hard to understand efforts put into such research endeavors except on the assumption that researchers thought it important to determine how location in allocentric space is encoded in the activity of place cells and to identify the sources from which places cells acquire information about location.

While this research attempted to show what content might be represented by place cell activity, it does not reveal whether the animal uses place cell activity as a repre-

¹³ [O'Keefe and Burgess \(1996\)](#) investigated how place fields might remap as a result of changes in the dimensions of the enclosure by comparing the responsiveness of the same neurons in enclosures that differed in the length of one or both walls—a small square, vertical rectangle (vertical wall extended), horizontal rectangle, and large square. This revealed that changing the length of the wall could cause place fields to expand, or sometimes split into two separate fields.

¹⁴ Both [Wills et al.](#) and [Leutgeb et al.](#) interpret their studies in light of the characterization of CA3 as constituting an attractor network. [Wills et al.](#) treat the abrupt transition in the place fields as indicative of an attractor network due to the recurrent connections in CA3 following the sparse coding imposed by DG. Although [Leutgeb et al.](#)'s results are in tension with a simple attractor network with a single global attractor, they construe the hysteresis effects as showing that more than a feedforward process is at work (perhaps recurrent networks with multiple attractors).

sentation of its location as it navigates its environment. There have been fewer investigations directed to this question. From the fact that performance on a navigation task remained above chance after a manipulation that created complete place cell remapping, Jeffery et al. (2003) reached a negative conclusion. They claimed that place cell remapping did not determine navigational behavior, suggesting that the information carried by place cells was not used as a representation in rat navigation. As Colgin et al. (2008) note, however, that the rat's behavior was significantly impaired after the manipulation (dropping from 91 % correct to 70 %) and propose that the retained success in performance may be due to additional neural processing that does not depend on the specific hippocampal map. Colgin et al. argue that place cells carry information that the rat uses in determining behavior.

6 Refining the account of how place cells represent places

While the line of research discussed above was addressing the question of how place cells acquired information about location, other research focused on other features of the activity of place cells that might serve as representations of the animal's location. The mere firing of place cells could provide at best a quite coarse-grained representation of an animal's spatial location—the animal could be anywhere in the place field. Initial investigations into how animals might represent space in a more precise manner focused on firing rate: the hypothesis was that if the rate of firing is distributed in a Gaussian manner around the center of the place field, the degree of reduction in firing rate below the maximal firing of the place cell would indicate how far the animal was from the center of the place field. However, the research on remapping described above indicated that the rate of firing carried non-spatial information. Even before that evidence came to light, however, Muller et al. (1987) argued against firing rate encoding fine-tuned spatial information. They contended that such a representational strategy would not work if a place cell has two or more fields, or has an oddly shaped field, and would produce systematic errors.

A different proposal for understanding how place cells can provide a more precise representation of location resulted from relating the generation of action potentials by place cells to EEG research that measured ongoing oscillations in electrical currents generated by ions moving across neural membranes. Oscillations of 6–12 Hz, labeled theta rhythms, had been identified in the hippocampus already in the 1930s¹⁵ and many researchers tried to link these oscillations with mental processes such as arousal after a strong stimulus, voluntary movement, and memory consolidation, but no consensus was achieved.¹⁶ The cells Ranck (see above) had identified as Θ -cells and that O'Keefe

¹⁵ The labeling of the frequency ranges of oscillations detected with EEG (or as local field potentials with implanted electrodes) as *alpha*, *beta*, reflects the order in which oscillations in a frequency range were discovered. The range labeled *theta* in the hippocampus extended higher (into the traditional alpha band) than the 4–8 Hz band associated with cortical theta waves.

¹⁶ See Buzsáki (2005) for a review and references. He concludes (p. 828): “Despite seven decades of hard work on rabbits, rats, mice, gerbils, guinea pigs, sheep, cats, dogs, old world monkeys, chimpanzees, and humans by outstanding colleagues, to date, there is no widely agreed term that would unequivocally describe behavioral correlate(s) of this prominent brain rhythm. By exclusion, the only firm message that

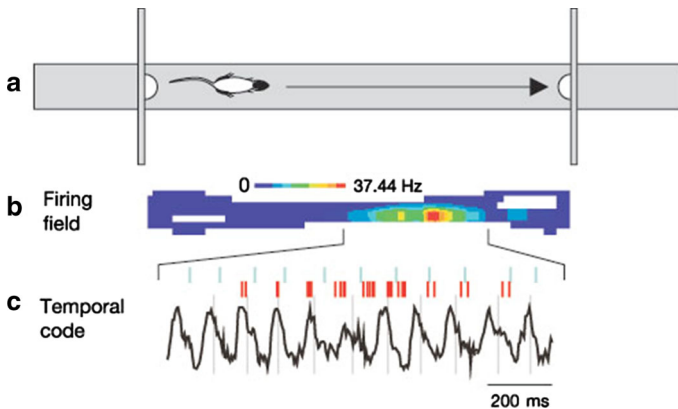


Fig. 4 Illustration of theta precession. As rat runs along the maze, it crosses the place field of a place cell (shown in the *middle*). The place cell spikes, shown in *red* at the *bottom*, precess against the underlying theta oscillation, firing first just after the peak and moving progressively earlier on subsequent theta cycles. Reprinted with permission of Nature Publishing Group from [Huxter et al. \(2003\)](#). (Color figure online)

had labeled as *displace cells* fired during periods when theta rhythms were evident. With regard to complex cells (O’Keefe’s place cells), Ranck found “no simple relation between the existence of a slow-wave theta rhythm and the firing of a complex spike cell.” He did note, however, that when a complex cell fired during a theta rhythm, the firing tended to be in phase with the theta rhythm. A number of researchers had pursued the relation between the behavior of complex cells and theta rhythms, but their focus was on finding a preferred theta phase for complex cell activity and to do so they averaged over multiple trials (see, e.g., [Buzsáki et al. 1983](#)). This concealed any systematic relation of individual place cell responses to theta.

O’Keefe and Recce (1993) pursued a different strategy, relating individual bursts from a single place cell in CA1 as the rat ran back and forth along a linear track (receiving rewards at each end) to the underlying theta rhythm. Typically on a single transit a rat would remain in the place field of an individual neuron for several theta oscillations, and the researchers noted that its spikes were regularly spaced at a slightly higher frequency than the prevailing theta rhythm. This ensured that the phase relation between the spikes and theta was not constant and O’Keefe and Recce determined that as the rat moved through the place field, successive spikes (or bursts of spikes) would occur earlier with respect to the theta cycle’s phase until the animal left the cell’s place field (by which point the spikes might have advanced nearly a full cycle). Thus, as shown in Fig. 4, the cell might emit a total of nine spike bursts in the course of only eight theta cycles. O’Keefe and Recce referred to this as *phase precession* and proposed that knowing how far the spikes had precessed against the prevailing theta rhythms provided a more accurate representation of the rat’s position than the place cell activity alone. This claim that was initially contested by several researchers but received com-

Footnote 16 continued

can be safely concluded from this brief summary is that in an immobile animal no theta is present, provided that no changes occur in the environment (and the animal is not ‘thinking’).”

elling support from [Jensen and Lisman \(2000\)](#) when they showed that by recording simultaneously from 38 hippocampal neurons they were able to reconstruct an animal's position on a linear track and that including phase information improved the accuracy of the reconstruction by 43 %. Thus, rather than firing rate, it is the time of spiking with respect to the theta cycle serves as the finer-grain representation of location.

The representational power of theta precession can be appreciated by focusing not on an individual neuron but on the activity over a population of neurons. At a given moment, the animal will be within the place fields of several place cells, some of which it is just entering and others that it has partly transgressed. The extent of precession of the different neurons within a theta cycle reveals its recent route. One might be skeptical, though, that such an esoteric code as provided by theta precession could really be a representation to the animal: who would be the consumer of such representations? In fact, however, such a temporal code is very useful for the hippocampus itself if we consider one of the challenges it faces in constructing a spatial map of a given locale: the various place cells must be integrated into a map by forming connections between cells representing adjacent locations (recall, representations for adjacent locations are not necessarily near each other but are connected functionally). This presumably involves long term potentiation (LTP), a process of enhancing the responsiveness of a neuron that produces an action potential after receiving input from a given neuron by increasing the number of NMDA receptors at synapses with the input neuron. LTP requires that the input neuron's spike occur within a very short time interval before the spike of the target neuron. Inputs from neurons spiking just before a target neuron in a given theta cycle (due to the animal having partly transgressed its place field) fit this requirement allowing these synapses to be strengthened. As the animal repeatedly explores the space, it will develop the connections between place cells that constitute the functional organization of a map ([Skaggs et al. 1996](#)). Moreover, once acquires, these associations can cause action potentials to occur in place cells that the animal has not yet reached and this can provide a representation of where the animal anticipates being in the future.

This research supports the claim I am advancing—that the object of the research is identifying the representational vehicles and their content. In this case the discovery of phase precession was a major factor prompting neuroscientists to extend their conception of the vehicles the brain could use beyond the firing rate of neurons that had long been the primary focus: the temporal specifics of the firing pattern could carry information independent of firing rate. Research on rate remapping discussed in the previous section already suggested that the firing rate of place cells may encode non-spatial information about the stimulus such as color or about the current behavior and goals of the animal. One can also easily envisage how coupling of information about such features with information about their spatial location can be useful—it provides a way of linking information about features of an environment and events happening there with the representation of the place in which they occurred.

7 Not all place cells do the same thing

Place cells have been identified in at least three different regions of the hippocampus, DG, CA3, and CA1, and so far I have not attended to any research into the differences

between place cells in different regions. The clear architectural differences between the regions, though, suggested to researchers there may be important differences. Prior to the discovery of grid cells (discussed in the following section), many researchers assumed that the construction of allocentric maps occurred within the hippocampus, possibly drawing upon the sparse coding exhibited in DG and the recurrent connections in CA3 before CA3 sent inputs onto CA1. Research on differences in the way remapping occurred in these different areas began to point to a more complex picture. [Leutgeb et al. \(2004\)](#) began by investigating how much the representations for different rooms overlapped in CA3 and CA1. They found little overlap in CA3 but substantial overlap in CA1. If different enclosures were employed in the different rooms, the overlap in CA1 diminished. They took this to indicate that CA3 is more involved in differentiating locations whereas CA1 is more involved in responding to similarities. They then introduced a novel room and investigated how quickly differentiated mappings occurred. The new room generated a distinct response immediately in CA3, although it took 20 min or more for new place fields to stabilize (likely a result of the recurrent connections within CA3). In contrast the representations of the new room in CA1 arose almost immediately and underwent little change. [Leutgeb et al.](#) interpreted this as indicating that CA1 must be relying on direct input from entorhinal cortex.

[Leutgeb et al.](#)'s research highlighted the direct pathway from EC to CA1, but there is also the indirect pathway from CA3 to CA1. Some researchers were curious as to how and when CA3 affects CA1. An interesting proposal, advanced by [Colgin et al. \(2009\)](#), is that the CA3 input is dominant when the response in CA3 indicates a very close match to an environment for which a map already exists. The input from CA3 to CA1 reinstates the previously learned map in CA1. When this is not the case, that is, the input is construed as new, then the direct pathway from EC to CA1 dominates, and CA1 develops a new response for the new input. This proposal is supported by findings about the temporal dynamics of the two pathways. Embedded within the slower theta oscillations found in LFPs are faster oscillations that fall within the gamma band (> 30 Hz). EC and DG appear to be generators respectively of faster (> 60 Hz) and slower (< 60 Hz) gamma oscillations and at different times the common target neurons in CA1 exhibit oscillations synchronized with the fast and slow gamma oscillations. The importance of this is that a post-synaptic cell is more likely to fire in response to an input from a presynaptic cell when their sub-threshold oscillations are synchronized.¹⁷ This indicates that CA1 switches between responding to EC and CA3 and one could hypothesize that switching itself would be governed by whether CA3 succeeded in settling into a previously learned attractor or not.

¹⁷ For a discussion of this and other roles oscillation and their synchronization play in neural computing, see [Buzsáki \(2005\)](#). In this paper I have focused on research analyzing the responses of individual neurons. This focus was a consequence of the fact that until recently researchers could only record from individual neurons. Increasingly, though, researchers are recording from populations of neurons. This enables them to focus on the interactive behavior. Regions such as CA3, with many recurrent connections, exhibit especially rich dynamic behavior, but in fact most brain regions exhibit a fair amount of recurrent connectivity between neurons, resulting in what Hebb referred to as cell assemblies. As [Buzsáki \(2010\)](#) argues, assemblies of neurons generating action potentials form dynamically in the brain through mechanisms that temporarily alter pre- and post-synaptic mechanisms. These mechanisms give rise to specific constellations of synaptic weights he refers to as *synapses*.




Fig. 5 The embedding of gamma oscillations within the theta cycle in CA1. The slow gamma oscillations occurring on the descending phase of theta oscillations shown in the two cycles on the left are synchronized with the slow gamma exhibited in CA 3 whereas the fast gamma oscillations occurring in the trough of the theta cycle shown in the right two cycles are synchronized with the fast gamma exhibited in EC. Reprinted by permission of the American Physiological Society from [Colgin and Moser \(2010\)](#)

Further support for this idea is provided by the discovery that fast and slow gamma oscillations occur at different phases of the theta cycle as recorded in CA1: fast gamma oscillations, synchronized with those in EC, occur in the trough of the theta cycle whereas the slow oscillations, synchronized with those in DG and CA3, occur on the descending portions of the theta wave (Fig. 5). [Hasselmo et al. \(2002\)](#) offered evidence that LTP does not occur at the peak or descending portions of the theta wave, when fast gamma oscillations are synchronized between CA1 and CA3. Rather, it occurs at the trough of theta, a period when CA1 is responding to EC but not CA3. This arrangement prevents LTP when CA3 recognizes an input as one already learned (preventing the overwriting of existing memories) but facilitates it when CA3 fails to recognize the input. [Colgin and Moser \(2010, p. 325\)](#) proposed that a further advantage of employing slow gamma for memory recall is that it is outside the optimal range for LTP (~ 25 ms); accordingly, “slow gamma oscillations may be fast enough to link distributed cells during memory retrieval yet also slow enough to avoid re-encoding of previously stored memories.” Thus, this proposal offers a mechanism for shutting off LTP when existing memories are being accessed.

One reason researchers put so much effort into determining what different neural processes represent is that they aim to understand how information is being processed in the overall mechanism. Knowing what different populations of cells represent serves as a guide to understanding the operations that transform one set of representations into another. The discovery of place cell remapping showed that place cell representations were changed over time. The differentiation of representation activities in CA3 and CA1 and the proposal for how communication between these regions and EC is gated constitutes a testable proposal for how new place cell representations are produced.

8 Situating place cells in a network of location-processing cells

So far I have focused on research on place cells within the hippocampus proper. For a place cell to represent an animal’s location, it must receive information in its inputs from which it can construct that representation. The inputs to the hippocampus must carry information about spatial location even if they do not represent that information in a format that can be used for other processing; the information processing involved

in constructing the representation extracts that information into such a format. As I noted above, for a long period after the discovery of place cells this was assumed to occur in the hippocampus itself. But the discovery of a variety of cell types in regions surrounding the hippocampus that are viewed as representing information relevant to creating and using maps has led researchers to expand their conception of what brain regions constitute the mechanism involved in spatial processing and to recharacterize their account of what place cells represent.

These discoveries of other cell types involved in spatial processing in areas adjacent to the hippocampus had roots in research in the last decades of the 20th century but researchers only advanced a cohesive account of the nature of the representations they provided and how they linked to place cell activity in the first decade of this century. I turn first to what have become known as grid cells, prominent especially in layer II of the MEC. In the 1980s and 1990s a number of researchers found evidence that cells in MEC, the principal source of inputs to DG and CA1, also seemed to be responsive to places. Quirk et al. (1992) reported that maximal firing of a MEC neuron was linked with a specific location, but found that the field that elicited responses seemed to be dispersed throughout the enclosure, not sharply delimited as with place cells. They characterized the coding as much more noisy than that in the hippocampus. Nonetheless, Quirk et al. performed studies comparable to those done with place cells to show that MEC cells were not responsive to particular behaviors and while their firing field shifted when the location of a cue card was rotated, they continued to fire in the absence of a cue card. Unlike hippocampal cells, the behavior of these cells was not affected by changing the shape of the enclosure. These findings led Quirk et al. to suggest that a role of the hippocampus was to enhance the signal received from the MEC (Muller et al. 1989, developed a simulation using back-propagation to show how place cell activity could be constructed from that of MEC cells).

The evidence that representation of space in MEC neurons seems to be at best noisy, however, reinforced the presumption of most researchers that the construction of an allocentric representation of space occurred within the hippocampus. A shortcoming of these earlier studies was that they had recorded from the ventrolateral, not the dorsolateral band of the MEC, and it is from the dorsolateral band that the projections to the relevant regions of the hippocampus originate. After Brun et al. (2002) showed that place cells continued to generate action potentials after their inputs along the indirect pathway from CA3 were severed, the Moser group recorded from the dorsolateral band. They found cells that had multiple “sharp and coherent place fields” that were “clearly delimited against the intervening background” (Fyhn et al. 2004, p. 1259). They ruled out the possibility that the spatial responses of MEC neurons resulted from the projections from the hippocampus back to the MEC by showing that they were maintained even after lesions to the hippocampus¹⁸ and concluded “The positional firing properties of the superficial dorsolateral band neurons suggest that spatial location is expressed accurately before signals enter the hippocampus” (p. 1260) and that “spatial signals are actively transformed into cohesive allocentric representations within

¹⁸ More recent research from the same group, however, has shown that MEC responses gradually diminish when the hippocampus is inhibited by GABA_A receptor agonist muscimol. I return to these studies below.

the entorhinal cortex itself” (p. 1261). (Fyhn et al. further found that by recording from several MEC cells they could reconstruct the path of the rat).

Having found evidence of location representation one step before the hippocampus, Fyhn et al. investigated whether such representations could be found at the next upstream brain region, the postrhinal cortex. There, however, they found no comparable cells. Moreover, others in the same research group showed that lesions to the dorsolateral band of the MEC resulted in impaired performance on navigation tasks in the Morris water maze (Steffenach et al. 2005). This suggested that not only did dorsolateral MEC neurons represent the animal’s location in space, but without it the animal could not perform the navigation tasks that had been associated with the hippocampus.

Fyhn et al. left unspecified whether there was any systematic relation between the different areas in space that activated a grid cell. There were suggestions that areas activating the same grid cell formed the vertices of equilateral triangles but the enclosure in which they had studied the rats was too small to support a stronger claim. The following year, recording from MEC while rats navigated in a larger enclosure that allowed them to develop a better perspective on the regions that elicited action potentials, Hafting et al. (2005) identified a distinctive layout of fields to which neurons in MEC responded: as shown in Fig. 6, three local fields define the vertices of an equilateral triangle and this pattern is repeated in a tessellating pattern across the whole local environment. Accordingly, around each local field the nearest six others formed a regular hexagon. Each neuron was thus associated with a grid of locations to which it would respond, and so these neurons were labeled *grid cells*. The grids for nearby neurons were offset from each other, but as in the hippocampus, nearby cells did not represent nearby regions of space.¹⁹

As had been the strategy following the discovery of place cells, the researchers followed up their discovery of grid cells by exploring their properties. In some ways the grid cells behave like place cells. For example, like place cells, grid cells exhibit phase precession (Hafting et al. 2008). One way grid cells differ from place cells is that the relationship between the fields represented by two grid cells remains constant across environments—when put in a new environment, two neurons will represent locations in the same spatial relationship to each other as in previous environments (Fyhn et al. 2007). Moreover, whereas the pattern of place cell responses to a new environment often takes time to develop, the grid structure was quickly manifest when a rat was introduced to a novel room (although it underwent refinement over subsequent exposure).

As I discussed above, one of the major strategies for investigating how place cells represented locations was to study remapping. Researchers have not found changes in the behavior of grid cells in response to environmental changes that, in the hippocampus, only result in rate remapping. Situations that generate global remapping in

¹⁹ The researchers, though, did find a distinct pattern across the entire MEC: the spacing of vertices and the size of the fields in which neurons responded expanded monotonically the further from the dorsal tip on the border with the perirhinal cortex (where the fields extended approximately 30 cm) to the ventral tip (3 m) (Brun et al. 2008). This mirrors the expansion of place fields as one moves along the dorsoventral axis of the hippocampus.




Fig. 6 Locations that elicit responses from a grid cell (*red dots*) as rat navigates path indicated by *black line*. Reprinted with permission of Annual Reviews from Moser et al. (2008). (Color figure online)

the hippocampus, however, also produce coherent shifts in the receptive fields of grid cells (Fyhn et al. 2007). Thus, when a cue card was the only visual cue on a circular enclosure and it was rotated, the grid structure rotated as well. But the cue card was not necessary and the grid structure remained largely in place (showing only a modest displacement of the vertices) in total darkness. Rather than undergoing a total remapping when, for example, a rectangular environment is stretched in one dimension, the grid is simply extended in that direction (Barry et al. 2007).

A noteworthy feature of the responses of place cells is that the cells that fire update as the animal moves. The animal may correct its location in light of environmental cues, but its first response is determined by its own movement (Redish et al. 2000). This requires that place cell activity be modulated by information about the animal's direction and speed of movement, a process referred to as *path integration*. Before grid cells were identified, researchers assumed path integration occurred in the hip-

pocampus itself. The fact that grid cells update in a similar fashion, a process that occurs when rats run in darkness and when they enter new environments, supported the idea that path integration occurs first in grid cells, which in turn drive the activity of place cells. The importance of MEC for path integration was demonstrated when Paron and Save (2004) showed that rat's with lesions in the entorhinal cortex could not perform a task that depends on path integration: returning to their starting location based on their self-motion.

A line of research related to that pursued on place cells in the 1980s and 1990s had identified yet another populations of cells, called *head direction cells*, that can provide the directional information. Ranck (1986) announced he had identified cells that responded to the absolute direction of the rat's head while recording from what he took to be the subiculum (when he completed the histology, he found the cells were actually in the postsubiculum). The full report that only appeared several years later (Taube et al. 1990a, b) described a Gaussian tuning curve with maximal response when the animal's head is facing in the target direction and tapered responses over 45° on either side of the direction of its maximum response. The researchers found that all compass directions are approximately equally represented in a population of head-direction cells. Muller, Ranck, and Taube (1996, p. 200) comment "The overall set of HD cells, therefore, acts as a compass that signals head orientation."²⁰ The characterization of these cells as constituting a compass is a bit misleading since Traube et al. had demonstrated that these cells responded to the relative direction in the local enclosure, not the geocentric direction in which the rat was facing. When a cue card in cylindrical enclosure was moved, the responses of these cells tracked the cue card (if it was rotated 90° , so was the direction to which the head-direction cells responded).²¹

A further feature of these cells, that they continued to track the direction the animal was facing even when all visual cues are removed or when the animals is moved from a familiar environment into a novel one,²² suggested that head direction cells secured information directly from vestibular cues and were in a position to supply it in the process of path integration. They typically would not do so perfectly, and when landmark information becomes available and indicates that the direction represented by the head direction cells has drifted from the correct one, the firing pattern is quickly corrected (Goodridge and Taube 1995). After their discovery in the postsubiculum, head direction cells were soon reported to be present in numerous other brain areas including the entorhinal cortex. Accordingly, the researchers who discovered grid cells in entorhinal cortex looked as well for head direction cells. Although they did not find them in layer II, where they had generated the strongest recording of grid

²⁰ As with place cells, there is not a topographical mapping of these cells onto directions. These cells only respond to direction in the horizontal plane and are largely unaffected by pitch or roll of the animal's head as well as its ongoing behavior. The pattern of activity is stable in a given environment over several days.

²¹ Using a task in which the animal had to make a correct choice to acquire a reward, Dudchenko and Taube (1997) showed that typically head direction cells rotated their firing fields, and then, and only then, did the animal perform the correct behavior. Other studies from the same laboratory have found more problematic relations between activity of head-direction cells and behavior (Golob et al. 2001).

²² When visual cues were removed, head direction and place cells both often rotated in unpredictable but correlated ways (Knierim et al. 1995).

cells, [Sargolini et al. \(2006\)](#) found head-direct and grid cells colocalized in the deeper layers (III–VI) of the entorhinal cortex. Many cells exhibited both head direction and grid field responses (grid \times head direction cells). They further found that firing of grid cells, head direction cells, and grid \times head direction cells is modulated by the animal's running speed, although where this information is represented and how it is integrated with head direction and grid location is not yet known.

The finding of both grid cells and head direction cells in MEC, as well as other discoveries,²³ provided the basis for a more comprehensive account of rodent navigation. On this account, path integration is performed in the deeper layers of MEC where information orientation, speed, and location information is integrated²⁴ and this information in turn is employed in developing the more abstract spatial representations found in the grid cells in layer II of MEC. Grid cells are viewed as providing the input to place cells in the hippocampus ([O'Keefe and Burgess 2005](#); [McNaughton et al. 2006](#)). This view radically alters the characterization of the role of the hippocampus in spatial processing.²⁵ Whereas the hippocampus had been viewed as the locus of allocentric spatial representation, the new view construed the area from which the hippocampus received its input as already providing an allocentric representation of space.

What then, is the contribution of place cells? Focusing on how grid and place cells differ in representing spaces, two things stand out that suggested an answer. One is that a given grid cells represents a variety of locations (all the grid locations that elicit activity in it), whereas place cells represent only one location (place field). The second

²³ The discovery of yet other kinds of cells also supports such an account. Independently, [Solstad et al. \(2008\)](#) and [Savelli et al. \(2008\)](#) discovered a type of cell in the dorsocaudal quarter of the MEC and adjacent parasubiculum that fires only when the rat is along one or sometimes several borders of an enclosure. Accordingly, they named these cells *border cells*. These cells typically fire along most of the length of a border, and if the enclosure is extended on the side of the border, their firing fields are likewise expanded. If a new wall is inserted parallel to the border, a new firing field emerges on the same side of that wall as the field at the enclosure boundary. When the walls are removed and the platform simply drops off at the previous boundary, the boundary cells continue to respond. The behavior of these cells is also not affected when one enclosure is substituted for another. Still another type of cell carrying spatial information, *boundary vector cells*, were initially predicted on the basis of computational models designed to account for the remapping of place fields as a result of stretching the enclosure in a given direction ([O'Keefe and Burgess 1996](#)). Individual cells were postulated to have receptive fields at different distances from the rat and the cells were hypothesized to fire whenever an enclosure boundary intersects that field, regardless of the direction in which the rat was moving. [Lever et al. \(2009\)](#) identified a population of such cells in the subiculum that fit this description. The firing of these cells was unaffected by such factors as the color or material of the boundary, and whether it was a wall or a drop. These cells do not remap in the wake of changes to enclosures that yield remapping of place cells or head direction cells. Recent research indicates that grid cells are also found in the presubiculum and the parasubiculum, where they are collocated with head direction and border cells ([Boccaro et al. 2010](#)).

²⁴ There is not space here to discuss research on how path integration might be performed. Most of this research has taken the form of developing computational models. Fundamentally, two modeling strategies have been pursued—oscillatory interference models in which the interference between multiple oscillators serves to track locations and attractor network models in which networks of interacting components create dynamically transformable attractors (see [Giocomo et al. 2011](#)).

²⁵ The fact that place cells typically carried information about more than spatial location was already clear in the remapping studies described above. For some researchers, especially those also influenced by research on human memory that viewed the hippocampus as fundamental to the encoding of episodic memories, this suggested a view in which the hippocampus was not an abstract map but served to relate spatial information to other content related to events the animal had experienced ([Eichenbaum et al. 1999](#)).

and perhaps more telling difference is that whereas grid cells do not seem to alter their behavior in light of cue information unless it is taken to indicate a new location, the firing rate of place cells represents contextual information. This suggested that perhaps the major activity of the hippocampus is not providing the spatial map, but linking additional information to the map. This provides a point of connection between animal research on the hippocampus and human research, which has emphasized its role in episodic memory. The environmental contexts that are represented are the sorts of things that one might remember about an event occurring at a given location.

As I have presented this view, the place cell representations are dependent on the representations provided by grid cells and head direction cells. However, as I noted at the outset, the outputs from the hippocampus project back to the entorhinal cortex and they may affect the representations provide in the MEC. For example, the location specified by grid cells may not cohere with the contextual features observed at a given location. In their conclusion, [Hafting et al. \(2005\)](#) suggest that this contextual information can, in such circumstances, feed back to reset the grid structure in MEC:

The contextual specificity of hippocampal representations suggests that during encoding, the hippocampus associates output from a generalized, path-integration-based coordinate system with landmarks or other features specific to the individual environment or context. Through backprojections to the superficial layers of the entorhinal cortex, associations stored in the hippocampus may reset the path integrator as errors accumulate during exploration. Anchoring the output of the path integrator to external reference points stored in the hippocampus or other cortical areas may enable alignment of entorhinal maps from one trial to the next, even when the points of departure are different (p. 805).

The feedback to the entorhinal cortex occurs both directly and as mediated by the subiculum, and arrives at layer V, from which there are extensive projections to layer II, the locus of many grid cells. Further support for this proposal was provided by the determination that the pattern of activity of grid cells destabilizes when the hippocampus is inactivated so that after several minutes the spatial structure is totally disrupted ([Hafting et al. 2008](#); [Bonnievie et al. 2013](#)). This suggests that grid representations are only stable when they support predictions, such as appropriate contextual cues, that place cells represent.

9 The consumption of spatial representations

So far I have focused on research seeking to identify representational vehicles and to determine what spatial information each carries. But spatial representations figure in information processing only if they enable the animal to navigate and solve problems involving spatial relations.²⁶ If this is to happen the spatial representations in the medial temporal cortex must affect processing in areas involved in motor planning.

²⁶ This point is nicely made by Moser and Moser “Finally, it is important to realize that the hippocampal–parahippocampal circuit only forms a representation. Representations can only influence navigation to the extent that information about the animal’s location is transferred to brain regions involved in planning and initiating movement” (2008, p. 1151).

Some research efforts have been devoted to this challenge. One of the major regions to which the parahippocampal and perirhinal cortexes project is the posterior parietal cortex (PPC). Multiple areas in the PCC are employed in transforming spatial information acquired through vision and other senses into the formats needed to control various bodily actions, such as eye movements, head movements, and grasping (Milner and Goodale 1995; Andersen and Buneo 2002). This processing often involves transforming a representation of egocentric space (e.g., specification of where something appears in a person's visual field) into a representation in the coordinate systems needed to direct the movement of the eyes, hands, or body.

Human neuroimaging studies support the hypothesis that the PPC figures in planning and executing navigational tasks (Rosenbaum et al. 2004; Spiers and Maguire 2007). Lesions to PPC in rodents result in deficits in both navigation based on cues (Kolb and Walkey 1987) and in path integration (Parron and Save 2004). Whitlock et al. (2008) interpret these deficits as resulting from the failure to translate allocentric spatial information as developed in the entorhinal and hippocampal areas into formats needed to direct movement. They also propose that the most likely route for linking these medial temporal areas to the PPC is via the retrosplinal cortex, which receives direct inputs from MEC, subiculum, and presubiculum and in turn has connections with PPC. To support this proposal they appeal to the fact that lesions to retrosplinal cortex result in disrupted spatial navigation and path integration.²⁷

A different strategy for demonstrating the role of the hippocampus and MEC in planning behavior has focused on other activities in these areas themselves. When we use external maps in planning routes we often interact with the map, tracing out possible routes before following the one we chose. There is some evidence that the hippocampus performs such operations. In addition to theta and gamma waves, a third rhythm is observed in the hippocampus: sharp waves appear either during sleep or when the animal is awake but still, as when resting after the completion of an activity or waiting to perform one (Buzsáki 1986). Diba and Buzsáki (2007) found that place cells that had fired in conjunction with theta rhythms while a rat was navigating fired in reverse order during the sharp waves that appeared when the animal stopped. A sequence of place cells also fired when the rat was preparing another run, which the researchers characterized as preplay (Fig. 7). Pastalkova et al. (2008) extended these findings in an experiment using a figure-eight maze in which rats were to run alternately in the two arms, but were first required to run on a running wheel. Place cells were even more active while the rats were on the running wheel than when actually running the maze. The pattern of place cell activity during running wheel behavior accurately predicted their subsequent running, including trials on which they chose the wrong

²⁷ Burke et al. (2005) addressed specifically the question of whether spatial differentiation was registered first in PPC and communicated to the hippocampal regions, or whether it was first developed in the hippocampal regions and then influenced PPC. Projections to the hippocampus originate in the deeper layers of PPC while the projections from the hippocampus are received in the superficial layers. The researchers showed, using an analysis of gene expression, that the differential representation of an enclosure in different rooms affected the superficial layers of PPC, not the deep layers, showing that the processing of location in the hippocampus regulate PPC behavior, not vice versa.




Fig. 7 **a** Firing of place cells as rat moves along a linear track. **b** Reverse replay of place cell activity as the rat rests at the end of the run. **c** Forward preplay of place cell firing as the rat prepares to run the course. Reprinted by permission of Annual Review from [Buckner \(2010\)](#)

arm.²⁸ While the mechanism that could initiate such preplay is not known, the results are invoked to show that hippocampal place cells are used in executing navigation tasks.

10 Conclusions

Philosophers such as Ramsey and Egan have expressed doubts as to whether the brain really contains representations, contending that attributing representations to different brain activities is simply a convenient and useful way for neuroscientists to describe what is happening. My strategy has been to demonstrate that the identification of representations in the brain is not, for the researchers I have discussed, a gloss in the presentation of research results but is the object of their inquiry. I have done that by showing that the research projects neuroscientists pursued on the role of the hippocampus and adjacent areas were directed at understanding what spatial information is being

²⁸ This study shared the limitation that the mazes used were far less complex than the routes rats navigate in the wild. Thus, despite its strong evidence for preplay, it does not conclusively answer the question of what role replay and preplay might subserve in real world navigation. One supportive indication came from Davidson, Kloosterman, and Wilson's (2009) demonstration that when rats are required to run on a straight track, ripple events get chained together for longer tracks. They suggested that this chaining might be mediated by reentrant processing between the hippocampus and entorhinal cortex, supporting the idea that replay during the SPW state might be a general mechanism employed in learning a navigational route.

represented in brain activity in different cell populations and how these representations figure in the information processing activities of the brain. The investment in determining just what and how brain processes represent suggests that the researchers are committed to representations as constituents of the information-processing mechanisms they study. They take their task to be is to discover the format employed by the representational system, to determine the content of these representations and how it is acquired it, and to figure out how these representations are then used in controlling behavior. Having built the commitment to representations, their content, and their use, into their project of understanding the mechanisms of cognition, neuroscientists would be perplexed by the suggestion that they were not fully committed to the brain activity really serves a representational function. It is not clear what their research has established if not what is represented in the rodent navigational mechanism and how these representations are used.

In the example I have developed O’Keefe’s identification of place cells in CA1 of the hippocampus set a research agenda. The goal was not just to understand neural processes in the hippocampus, although that had been the focus of inquiry in an earlier period and led to the discovery of LTP (Craver 2003). Rather, the researchers were concerned with explaining the phenomenon of spatial navigation which Tolman had argued required a map-like representation. In the vocabulary Bechtel and Richardson (1993)/2010 employed, the discovery of place cells constituted a *direct localization* of a map that provided an allocentric representation of space. Direct localization serves simply to locate where the phenomenon is produced—it doesn’t explain how it is generated by a mechanism. That requires either further decomposing the localized component into multiple parts or discovering that there are other parts in addition to that initially discovered. In the context of an information processing mechanism, the discovery of other relevant parts within a given part or outside it is often guided by investigating parts that seem to represent additional information that is needed in solving the problem. Both strategies were employed in research on spatial navigation—identifying differences in place cell behavior in different parts of the hippocampus and identifying other cell types beyond the hippocampus that interacted in representing space allocentrically and allow animals to use those representations to navigate. This inquiry involved determining what information place cells actually represented, what features of place cell activity served as the vehicle for representing particular kinds of information, and how they functioned to link information about places with information about context. In addition, they were interested in determining how these neurons acquired information about the place the animal was and how this information was subsequently used in a larger network of brain regions involved in spatial processing. These discoveries result in what Bechtel and Richardson called *complex localizations* in which multiple components contribute to the phenomenon sequentially, and sometimes in *integrated systems* if these parts interact with each other in multiple ways. An information processing mechanism is such an integrated system in which parts with different representational content interact.

Identifying representations is, I have argued, an important aspect of understanding the operation of information-processing mechanisms that are viewed as serving as control systems. In order to control something, a controller needs information about the plant that is being controlled or what the plant is interacting with. Neuroscientists

view many of the mechanisms they investigate as performing control functions and so treat them as processing information. This leads them to investigate what neural processes serve as representation and the contents they represent. Quite naturally, then, they engage in research directed at how neural mechanisms they identify as representations carry content about features of the environment and how operations performed on these representations figure in the control of behavior. Thus, as I have described, research on place cells, grid cells, and head direction cells has been directed at understanding how each represents spatial information in the mechanisms taken to be involved in processing spatial information of animal brains.

My focus has been on whether neuroscientists such as those investigating the role of populations of cells in the hippocampus and related areas are committed to these cells representing information. By examining the research they do, I have argued that they are committed to representations and that their talk of representations cannot be treated as mere glosses. I have not addressed the question of whether they are correct in their assumption that there are representations in the brain. As I noted in the introduction, there are research programs that seek to explain the role of the brain in producing behavior that do not traffic in representations. The fate of a research program in the face of such a clash is typically settled by whether its practitioners view themselves as making progress, and those pursuing information-processing explanations that appeal to representations certainly see themselves as making progress. In the future they may conclude that their approach has stalled and abandon it. If they do, it is not clear what results of their research will survive since what they are trying to develop is an account that identifies representations and their content. Whether researchers eventually abandon the quest for explaining brain activity in terms of information-processing mechanisms, my contention is that as long as they pursue this endeavor, characterizing those processes in representational terms is not a mere interpretive gloss but rather foundational to their inquiry.

Acknowledgments I conducted the research for this paper while a Fellow at the Institute for Advanced Studies at Hebrew University and their support is gratefully acknowledged. I thank members of the Mind and Brain Research Group at the Institute for productive discussions. I also received valuable feedback from audiences at the Center for Theoretical Neuroscience at the University of Waterloo, the Southern Society for Philosophy and Psychology, the University of Amsterdam, and the Rotman Institute at the University of Western Ontario.

References

- Abrahamsen, A., & Bechtel, W. (2011). From reactive to endogenously active dynamical conceptions of the brain. In K. Plaisance & T. Reydon (Eds.), *Philosophy of behavioral biology* (pp. 329–366). New York: Springer.
- Andersen, R. A., & Buneo, C. A. (2002). Intentional maps in posterior parietal cortex. *Annual Review of Neuroscience*, 25, 189–220.
- Barry, C., Hayman, R., Burgess, N., & Jeffery, K. J. (2007). Experience-dependent rescaling of entorhinal grids. *Nature Neuroscience*, 10, 682–684.
- Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. London: Routledge.
- Bechtel, W. (2011). Representing time of day in circadian clocks. In A. Newen, A. Bartels, & E.-M. Jung (Eds.), *Knowledge and representation*. Palo Alto, CA: CSLI Publications.

- Bechtel, W., & Abrahamsen, A. (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences*, *36*, 421–441.
- Bechtel, W., & McCauley, R. N. (1999). Heuristic identity theory (or back to the future): The mind-body problem against the background of research strategies in cognitive neuroscience. In M. Hahn & S. C. Stoness (Eds.), *Proceedings of the 21st annual meeting of the cognitive science society* (pp. 67–72). Mahwah, NJ: Lawrence Erlbaum Associates.
- Bechtel, W., & Richardson, R. C. (1993/2010). *Discovering complexity: Decomposition and localization as strategies in scientific research*. Cambridge, MA: MIT Press. 1993 edition published by Princeton University Press.
- Boccaro, C. N., Sargolini, F., Thoresen, V. H., Solstad, T., Witter, M. P., Moser, E. I., et al. (2010). Grid cells in pre- and parasubiculum. *Nature Neuroscience*, *13*, 987–994.
- Bonnevie, T., Dunn, B., Fyhn, M., Hafting, T., Derdikman, D., Kubie, J. L., et al. (2013). Grid cells require excitatory drive from the hippocampus. *Nature Neuroscience*, *16*, 309–317.
- Bostock, E., Muller, R. U., & Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, *1*, 193–205.
- Brun, V. H., Otnæss, M. K., Molden, S., Steffenach, H.-A., Witter, M. P., Moser, M.-B., et al. (2002). Place cells and place recognition maintained by direct entorhinal–hippocampal circuitry. *Science*, *296*, 2243–2246.
- Brun, V. H., Solstad, T., Kjelstrup, K. B., Fyhn, M., Witter, M. P., Moser, E. I., et al. (2008). Progressive increase in grid scale from dorsal to ventral medial entorhinal cortex. *Hippocampus*, *18*, 1200–1212.
- Buckner, R. L. (2010). The role of the hippocampus in prediction and imagination. *Annual Review of Psychology*, *61*, 27–48.
- Burgess, N., & O’Keefe, J. (2011). Models of place and grid cell firing and theta rhythmicity. *Current Opinion in Neurobiology*, *21*, 734–744.
- Burke, S. N., Chawla, M. K., Penner, M. R., Crowell, B. E., Worley, P. F., Barnes, C. A., et al. (2005). Differential encoding of behavior and spatial context in deep and superficial layers of the neocortex. *Neuron*, *45*, 667–674.
- Buzsáki, G. (1986). Hippocampal sharp waves: Their origin and significance. *Brain Research*, *398*, 242–252.
- Buzsáki, G. (2005). Theta rhythm of navigation: Link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus*, *15*, 827–840.
- Buzsáki, G. (2010). Neural syntax: Cell assemblies, synapse ensembles, and readers. *Neuron*, *68*, 362–385.
- Buzsáki, G., Lai-Wo, S. L., & Vanderwolf, C. H. (1983). Cellular bases of hippocampal EEG in the behaving rat. *Brain Research Reviews*, *6*, 139–171.
- Chemero, A. (2000). Anti-representationalism and the dynamical stance. *Philosophy of Science*, *67*, 625–647.
- Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge, MA: MIT Press.
- Chomsky, N. (1995). Language and nature. *Mind*, *104*, 1–61.
- Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., et al. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, *462*, 353–357.
- Colgin, L. L., & Moser, E. I. (2010). Gamma oscillations in the hippocampus. *Physiology*, *25*, 319–329.
- Colgin, L. L., Moser, E. I., & Moser, M.-B. (2008). Understanding memory through hippocampal remapping. *Trends in Neurosciences*, *31*, 469–477.
- Craver, C. F. (2003). The making of a memory mechanism. *Journal of the History of Biology*, *36*, 153–195.
- Cummins, R. (1989). *Meaning and mental representation*. Cambridge, MA: MIT Press.
- Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron*, *63*, 497–507.
- Diba, K., & Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, *10*, 1241–1242.
- Dretske, F. I. (1988). *Explaining behavior: Reasons in a world of causes*. Cambridge, MA: MIT Press.
- Dudchenko, P. A., & Taube, J. S. (1997). Correlation between head direction cell activity and spatial behavior on a radial arm maze. *Behavioral Neuroscience*, *111*, 3–19.
- Egan, F. (2010). Computational models: A modest role for content. *Studies in History and Philosophy of Science*, *41*, 253–259.
- Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., & Tanila, H. (1999). The hippocampus, memory, and place cells: Is it spatial memory or a memory space? *Neuron*, *23*, 209–226.
- Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., & Moser, E. I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex. *Nature*, *446*, 190–194.

- Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., & Moser, M.-B. (2004). Spatial representation in the entorhinal cortex. *Science*, *305*, 1258–1264.
- Garzón, F., & Rodríguez, Á. (2009). Where is cognitive science heading? *Minds and Machines*, *19*, 301–318.
- Garzon, P. C., & Keijzer, F. (2011). Plants: Adaptive behavior, root-brains, and minimal cognition. *Adaptive Behavior*, *19*, 155–171.
- Giocomo, L. M., Moser, M. B., & Moser, E. I. (2011). Computational models of grid cells. *Neuron*, *71*, 589–603.
- Golob, E. J., Stackman, J. R. W., Wong, A. C., & Taube, J. S. (2001). Correlation between head direction cell activity and spatial behavior on a radial arm maze. *Behavioral Neuroscience*, *115*, 285–304.
- Goodridge, J. P., & Taube, J. S. (1995). Preferential use of the landmark navigational system by head direction cells in rats. *Behavioral Neuroscience*, *109*, 49–61.
- Hafting, T., Fyhn, M., Bonnevie, T., Moser, M.-B., & Moser, E. I. (2008). Hippocampus-independent phase precession in entorhinal grid cells. *Nature*, *453*, 1248–1252.
- Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*, 801–806.
- Haselager, P., de Groot, A., & van Rappard, H. (2003). Representationalism versus anti-representationalism: A debate for the sake of appearance. *Philosophical Psychology*, *16*, 5–23.
- Hasselmo, M. E., Bodelón, C., & Wyble, B. P. (2002). A proposed function for hippocampal theta rhythm: Separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Computation*, *14*, 793–817.
- Huxter, J., Burgess, N., & O’Keefe, J. A. (2003). Independent rate and temporal coding in hippocampal pyramidal cells. *Nature*, *425*, 828–832.
- Jeffery, K. J., Gilbert, A., Burton, S., & Strudwick, A. (2003). Preserved performance in a hippocampal-dependent spatial task despite complete place cell remapping. *Hippocampus*, *13*, 175–189.
- Jensen, O., & Lisman, J. E. (2000). Position reconstruction from an ensemble of hippocampal place cells: Contribution of theta phase coding. *Journal of Neurophysiology*, *83*, 2602–2609.
- Jirsa, V. K., Fuchs, A., & Kelso, J. A. S. (1998). Connecting cortical and behavioral dynamics: Bimanual coordination. *Neural Computing*, *10*, 2019–2045.
- Kaplan, D. M., & Bechtel, W. (2011). Dynamical models: An alternative or complement to mechanistic explanations? *Topics in Cognitive Science*, *3*, 438–444.
- Kelso, J. A. S. (1995). *Dynamic patterns: The self organization of brain and behavior*. Cambridge, MA: MIT Press.
- Knierim, J., Kudrimoti, H., & McNaughton, B. (1995). Place cells, head direction cells, and the learning of landmark stability. *The Journal of Neuroscience*, *15*, 1648–1659.
- Kolb, B., & Walkley, J. (1987). Behavioural and anatomical studies of the posterior parietal cortex in the rat. *Behavioural Brain Research*, *23*, 127–145.
- Kuhn, T. S. (1970). *The structure of scientific revolutions* (2nd ed.). Chicago: University of Chicago Press.
- Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge* (pp. 91–196). Cambridge: Cambridge University Press.
- Leutgeb, J. K., Leutgeb, S., Treves, A., Meyer, R., Barnes, C. A., McNaughton, B. L., et al. (2005a). Progressive transformation of hippocampal neuronal representations in “Morphed” environments. *Neuron*, *48*, 345–358.
- Leutgeb, S., Leutgeb, J. K., Barnes, C. A., Moser, E. I., McNaughton, B. L., & Moser, M.-B. (2005b). Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science*, *309*, 619–623.
- Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M.-B., & Moser, E. I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science*, *305*, 1295–1298.
- Lever, C., Burton, S., Jeewajee, A., O’Keefe, J., & Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *The Journal of Neuroscience*, *29*, 9771–9777.
- Lever, C., Wills, T., Cacucci, F., Burgess, N., & O’Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature*, *416*, 90–94.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, *67*, 1–25.
- Markus, E. J., Qin, Y., Leonard, B., Skaggs, W. E., McNaughton, B. I., & Barnes, C. A. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *Journal of Neuroscience*, *15*, 7079–7094.

- Marr, D. C. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London B*, 262, 23–81.
- Marr, D. C. (1982). *Vision: A computation investigation into the human representational system and processing of visual information*. San Francisco: Freeman.
- McCauley, R. N., & Bechtel, W. (2001). Explanatory pluralism and the heuristic identity theory. *Theory and Psychology*, 11, 736–760.
- McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.-B. (2006). Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience*, 7, 663–678.
- Miller, V. M., & Best, P. J. (1980). Spatial correlates of hippocampal unit activity are altered by lesions of the fornix and entorhinal cortex. *Brain Research*, 194, 311–323.
- Milner, A. D., & Goodale, M. G. (1995). *The visual brain in action*. Oxford: Oxford University Press.
- Moita, M. A. P., Rosis, S., Zhou, Y., LeDoux, J. E., & Blair, H. T. (2004). Putting fear in its place: Remapping of hippocampal place cells during fear conditioning. *The Journal of Neuroscience*, 24, 7015–7023.
- Morgan, A. (2014). Representations gone mental. *Synthese*, 191, 213–244.
- Moser, E. I., & Moser, M.-B. (2008). A metric for space. *Hippocampus*, 18, 1142–1156.
- Moser, E. I., Kropff, E., & Moser, M.-B. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annual Review of Neuroscience*, 31, 69–89.
- Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience*, 7, 1951–1968.
- Muller, R. U., Kubie, J. L., Bostock, E., Traube, J. S., & Quirk, G. (1991). Spatial firing correlates of neurons in the hippocampal formation of freely moving rats. In J. Paillard (Ed.), *Brain and space* (pp. 296–333). Oxford: Oxford University Press.
- Muller, R. U., Kubie, J. L., & Ranck, J. B. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *The Journal of Neuroscience*, 7, 1935–1950.
- Muller, R. U., Quirk, G. J., & Kubie, J. L. (1989). Back-propagation calculations of hippocampal place cell firing fields. *Society for Neuroscience Abstracts*, 15, 403.
- Muller, R. U., Ranck, J. B., & Taube, J. S. (1996). Head direction cells: Properties and functional significance. *Current Opinion in Neurobiology*, 6, 196–206.
- Navratilova, Z., Giocomo, L. M., Fellous, J.-M., Hasselmo, M. E., & McNaughton, B. L. (2012). Phase precession and variable spatial scaling in a periodic attractor map model of medial entorhinal grid cells with realistic after-spike dynamics. *Hippocampus*, 22, 772–789.
- O'Keefe, J. A. (1976). Place units in the hippocampus of the freely moving rat. *Experimental Neurology*, 51, 78–109.
- O'Keefe, J. A., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381, 425–428.
- O'Keefe, J. A., & Burgess, N. (2005). Dual phase and rate coding in hippocampal place cells: Theoretical significance and relationship to entorhinal grid cells. *Hippocampus*, 15, 853–866.
- O'Keefe, J. A., & Conway, D. H. (1978). Hippocampal place units in the freely moving rat: Why they fire where they fire. *Experimental Brain Research*, 31, 573–590.
- O'Keefe, J. A., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely moving rat. *Brain Research*, 34, 171–175.
- O'Keefe, J. A., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Oxford University Press.
- O'Keefe, J. A., & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3, 317–330.
- O'Keefe, J. A., & Speakman, A. (1987). Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research*, 68, 1–27.
- Parron, C., & Save, E. (2004). Evidence for entorhinal and parietal cortices involvement in path integration in the rat. *Experimental Brain Research*, 159, 349–359.
- Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, 321, 1322–1327.
- Quirk, G., Muller, R., & Kubie, J. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *The Journal of Neuroscience*, 10, 2008–2017.
- Quirk, G., Muller, R., Kubie, J., & Ranck, J. (1992). The positional firing properties of medial entorhinal neurons: Description and comparison with hippocampal place cells. *The Journal of Neuroscience*, 12, 1945–1963.
- Ramsey, W. (2007). *Representation reconsidered*. Cambridge, NY: Cambridge University Press.

- Ranck, J. B. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. Part I. Behavioral correlates and firing repertoires. *Experimental Neurology*, *41*, 462–531.
- Ranck, J. B. (1986). Head direction cells in the deep cell layer of dorsal presubiculum in freely moving rats. In G. Buzsáki & C. H. Vanderwolf (Eds.), *Electrical activity of the archicortex* (pp. 217–220). Budapest: Academia Kiado.
- Redish, A. D., Rosenzweig, E. S., Bohanick, J. D., McNaughton, B. L., & Barnes, C. A. (2000). Dynamics of hippocampal ensemble activity realignment: Time versus space. *The Journal of Neuroscience*, *20*, 9298–9309.
- Rosenbaum, R. S., Ziegler, M., Winocur, G., Grady, C. L., & Moscovitch, M. (2004). “I have often walked down this street before”: fMRI Studies on the hippocampus and other structures during mental navigation of an old environment. *Hippocampus*, *14*, 826–835.
- Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B. L., Witter, M. P., Moser, M.-B., et al. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, *312*, 758–762.
- Savelli, F., Yoganarasimha, D., & Knierim, J. J. (2008). Influence of boundary removal on the spatial representations of the medial entorhinal cortex. *Hippocampus*, *18*, 1270–1282.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery, and Psychiatry*, *20*, 11–21.
- Shagrir, O. (2010). Marr on computational-level theories. *Philosophy of Science*, *77*, 477–500.
- Shagrir, O. (2012). Structural representations and the brain. *The British Journal for the Philosophy of Science*, *63*(3), 519–545.
- Shagrir, O., & Bechtel, W. (2011). Marr’s computational-level theories and delineating phenomena. In D. M. Kaplan (Ed.), *Integrating psychology and neuroscience: Prospects and problems*. Oxford: Oxford University Press.
- Shapiro, J. A. (2007). Bacteria are small but not stupid: Cognition, natural genetic engineering and socio-bacteriology. *Studies in History and Philosophy of Science Part C*, *38*, 807–819.
- Skaggs, W. E., & McNaughton, B. L. (1998). Spatial firing properties of hippocampal CA1 populations in an environment containing two visually identical regions. *The Journal of Neuroscience*, *18*, 8455–8466.
- Skaggs, W. E., McNaughton, B. L., Wilson, M. A., & Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, *6*, 149–172.
- Solstad, T., Boccara, C. N., Kropff, E., Moser, M.-B., & Moser, E. I. (2008). Representation of geometric borders in the entorhinal cortex. *Science*, *322*, 1865–1868.
- Spiers, H. J., & Maguire, E. A. (2007). The neuroscience of remote spatial memory: A tale of two cities. *Neuroscience*, *149*, 7–27.
- Steffenach, H. A., Witter, M., Moser, M. B., & Moser, E. I. (2005). Spatial memory in the rat requires the dorsolateral band of the entorhinal cortex. *Neuron*, *45*, 301–313.
- Swoyer, C. (1991). Structural representation and surrogative reasoning. *Synthese*, *87*, 449–508.
- Taube, J. S., Muller, R. U., & Ranck, J. B. (1990a). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *The Journal of Neuroscience*, *10*, 420–435.
- Taube, J. S., Muller, R. U., & Ranck, J. B. (1990b). Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *The Journal of Neuroscience*, *10*, 436–447.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*, 189–208.
- van Gelder, T., & Port, R. (1995). It’s about time: An overview of the dynamical approach to cognition. In R. Port & T. van Gelder (Eds.), *It’s about time*. Cambridge, MA: MIT Press.
- Whitlock, J. R., Sutherland, R. J., Witter, M. P., Moser, M.-B., & Moser, E. I. (2008). Navigating from hippocampus to parietal cortex. *Proceedings of the National Academy of Sciences*, *105*, 14755–14762.
- Wills, T. J., Lever, C., Cacucci, F., Burgess, N., & O’Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. *Science*, *308*, 873–876.
- Zinyuk, L., Kubik, S., Kaminsky, Y., Fenton, A. A., & Bures, J. (2000). Understanding hippocampal activity by using purposeful behavior: Place navigation induces place cell discharge in both task-relevant and task-irrelevant spatial reference frames. *Proceedings of the National Academy of Sciences*, *97*, 3771–3776.