3

Navigating in a Three-Dimensional World

Kathryn J. Jeffery

Abstract

A central theme in the study of animal navigation has been the extent to which such navigation requires the formation of an internal representation of space, the so-called "cognitive map." Although its properties remain disputed, it is now generally accepted that a map-like representation exists in the brain, and neurobiological studies, conducted in tandem with behavioral investigations, have done much to elucidate the neural substrate of the map as it operates in two dimensions. However, to date little is known about how the map encodes real-world, three-dimensional space. Using recent neurobiological and behavioral findings, this issue is explored here. It is argued that the navigational problems in three dimensions are qualitatively as well as quantitatively different from those in two dimensions, and evidence suggests that, perhaps for this reason, horizontal and vertical space are processed separately in the vertebrate brain. The possible adaptive consequences of such an anisotropic encoding scheme are outlined.

Introduction

Navigation has been one of the most intensively studied aspects of animal cognition in the past century because it is a fundamental competence, and because it requires engagement of a number of processes that are also relevant to cognition more generally, such as memory and planning. A central theme in spatial cognition research has been whether there exists an internal representation of allocentric (world-centered) space, the so-called "cognitive map" (Tolman 1948). With the advent of single neuron studies in freely behaving animals, evidence has steadily accrued that this is undoubtedly the case: there is an internal representation of space in the mammalian brain, located in the hippocampal formation and surrounding regions. This representation has many of the properties to be expected of a map, and thus the hippocampal spatial representation is often referred to as the "cognitive map."

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

Although self-localization, navigation, and the cognitive map have been well studied in two dimensions (2D), the world actually exists in three dimensions (3D). Even animals that are constrained to move over a surface, because they cannot swim or fly, frequently travel over undulating terrain, which complicates distance-measuring (odometry). In addition, many animals, including our evolutionarily distant ancestors (fish) and our close relatives (monkeys and apes), travel freely through volumetric space, as when swimming, gliding or flying, climbing through burrow systems, or swinging through trees. While we humans lost this capacity temporarily, we have recently regained it with the development of undersea, air, and space exploration and, most recently, virtual reality worlds in which we, or our avatars, can "move" freely in all dimensions. Given that so many of our evolutionary ancestors moved through volumetric space, the brain is likely to have evolved a mechanism long ago for representing space in all three dimensions, and not just two of them. However, the question of how this might be achieved has remained unexplored. In this chapter, I address the question of whether the cognitive map is three dimensional, like the world is, or whether it is two dimensional and encodes only horizontal space.

Surprisingly, the question of the 3D structure of the cognitive map has received relatively little attention to date, partly because of the technical difficulties of collecting behavioral and neural data in three dimensions, and partly because it has been assumed that encoding of the vertical dimension is probably just like encoding of the other two. However, encoding three dimensions is fundamentally different from encoding two, qualitatively as well as quantitatively. Recent evidence suggests that the vertical and horizontal components of space may be represented differently, and perhaps separately, in the vertebrate brain; that is, the cognitive map is *anisotropic*. This chapter reviews the neural basis of the cognitive map and examines the question of whether it encodes 3D space, and if so, whether this encoding is indeed anisotropic, and why this might be.

The Cognitive Map in Two Dimensions

The core neural substrate of the mammalian cognitive map is the place cell system, discovered by O'Keefe and colleagues in the hippocampus of freely moving rats in the early 1970s (O'Keefe and Dostrovsky 1971). In this experimental paradigm, recordings are made of the activity of single neurons detected by chronically implanted, in-dwelling electrodes, while animals explore an environment (Figure 3.1a). Place cells are hippocampal neurons that become active in spatially localized regions of the environment, known as place fields, and are thought to encode collectively a representation of the environment and the animal's location within it. Following the discovery of place cells, neurons in surrounding brain regions were discovered that respond to head direction

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

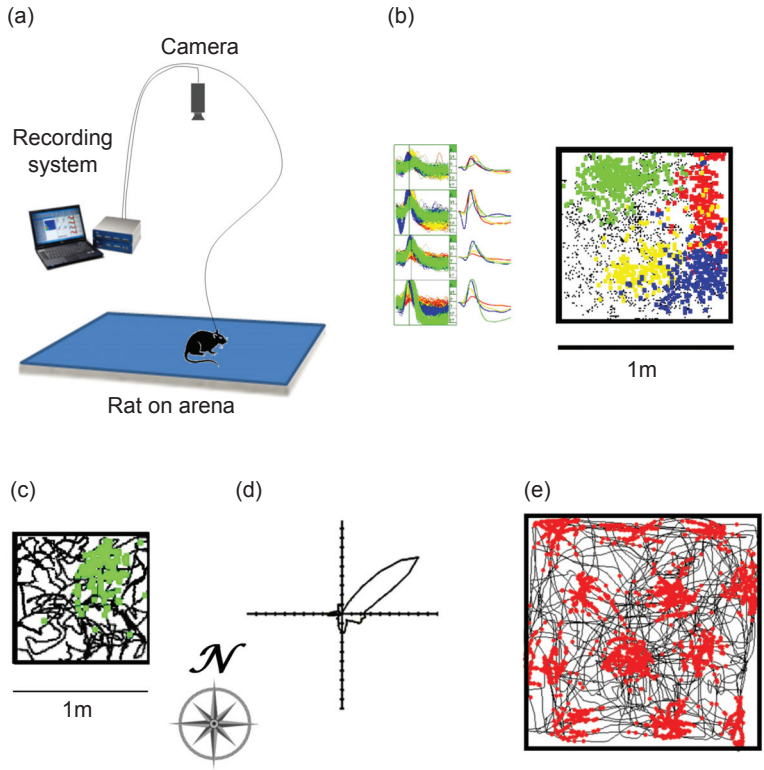

Figure 3.1 Recording spatially sensitive neurons in behaving animals. (a) Typical experimental setup. A rat explores an arena, foraging for food, while recordings of single neurons are made. The signals (action potentials emitted by single neurons) from chronically implanted electrodes are conveyed via a thin recording cable to the recording system. Meanwhile, an overhead camera tracks the position of the animal. (b) An example of a set of single neuron data. The waveforms on the left are action potentials, whose signals were collected simultaneously on four electrodes. Because the potentials from different neurons have slightly different shapes, it is possible to determine which signals came from which cell (color-coded on the figure). The pictures on the right show the action potentials represented as dots, with each dot color coded as for the waveforms and each placed on the location of the rat at the moment the cell fired. These neurons, which were hippocampal place cells, concentrated their action potentials ("spikes") in a spatially localized region of the environment. Note that different neurons concentrated their spikes in different places, so the environment as a whole is represented by the neuronal population. (c–e) Comparison of data from (c) a place cell, (d) a head direction cell, and (e) a grid cell. An arbitrary north is depicted by the compass rose. (c) The place cell has a single firing location ("place field") as shown by the concentration of spikes (green spots) in one part of the arena. (d) The head direction cell's activity is shown as a polar plot of firing rate against head direction. Note that firing rate increases markedly when the rat's head is facing north–east. (e) The grid cell, like the place cell, also lays its action potentials (red dots) down in focal parts of the environment, but unlike the place cell it has multiple place fields and these are arranged in a regular array, suggesting they may function like hexagonal "graph paper."

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

(for a review, see Taube 2007), hence their name "head direction cells." These cells are thought to function in, among other things, the orientation of the place cell map. Most recently, neurons called "grid cells" have been identified in the entorhinal cortex, upstream of the hippocampus (Hafting et al. 2005), whose function seems to be to encode directionally tagged distances.

The findings of place cells, head direction (HD) cells, and grid cells suggest that the brain is equipped with the essential tools with which to construct and use a map: a compass (HD cells), an odometer (grid cells), and the map itself (place cells). Anatomically speaking, the main direction of information flow is HD cells → entorhinal grid cells → hippocampal place cells, with return connections from hippocampus to the deep layers of entorhinal cortex. Study of these cells, and how they interact, can reveal important clues as to how animals encode space. What follows is a brief discussion of their properties as documented in 2D space, followed by a look at how these neurons may help us understand the encoding of 3D space.

The Place Cell Map in Two Dimensions

We begin with the place cells and the question of how the place cell "map" appears to be constructed from incoming sensory inputs. An example of the remarkable, spatially localized activity—the place field—of a place cell is shown in Figure 3.1b. Early studies of place cells rapidly determined that they are anchored to physical—often visual—features of the environment, because translation or rotation of such features caused place fields to move accordingly (O'Keefe and Dostrovsky 1971; Muller and Kubie 1987). It seems that for place cells, the boundaries of an environment are especially important for anchoring. This was confirmed by a study of O'Keefe and Burgess (1996), in which an environment was transformed by squashing or stretching it; the place cell map also squashed or stretched, although by a lesser amount, suggesting that the place cell map may be attached to the environment by its "edges."

What about place fields that are not *at* the boundaries of the environment, and thus lack its direct sensory (e.g., tactile) inputs? How do they manage to localize their activity so precisely? How does a given cell determine when the rat is at the correct distance from the walls and that it should start firing? A priori two broad possibilities arise, which can be classed as nonmetric versus metric. According to the nonmetric view, place cells are driven by inputs that may correlate with distances and directions but do not directly encode them. For example, if a cell responds to a set of stimuli, such as a visual snapshot (i.e., the exact view from a particular vantage point in the box), and cannot fire in the absence of these specific stimuli, this would constitute nonmetric spatial encoding: although the firing is spatially localized, no distances or directions are explicitly represented. One could not do trigonometry, for example, with such an encoding scheme. On the other hand, if the cell also fires in that same location in the dark when the animal can be shown (by interventional

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

manipulations) to be using only self-motion cues (also known as *path integration* cues), then the cell must be using metric encoding. This is because the only way that visual and self-motion cues resemble each other is that they convey the same information about position; that is, they are *isomorphic*, or informationally equivalent (Gallistel 1990) in the metric domain. Both kinds of sensory information, visual and path integration, are thus converted to a signal that "means" distance (or direction).

The metric view, by contrast, supposes that the cells receive information that explicitly encodes metric information (i.e., distances and directions with respect to fixed features of the environment). Thus, a place cell whose field is destined to be somewhere near the center of the environment receives explicit information about distance and direction from the anchoring boundaries. Initial evidence suggested that place cell encoding in the horizontal plane is probably metric, because the cells indeed show the same patterns of activity both when visual cues are available, in which case they control the activity of the cells (Muller and Kubie 1987), or unavailable, in which case the cells use path integration (Jeffery 2007). Furthermore, place fields are often positioned at a fixed distance from a barrier that is located *behind* the animal (Gothard et al. 1996), again suggesting a distance calculation based on self-motion. The source of direction information is thought to be from a system known as the HD system, while the distance information may come from odometric processes that are combined with the directional signal in a set of neurons known as the grid cells. Head direction and grid cell systems are discussed further below. A more comprehensive description of how place, head direction, and grid cells cooperate to form the place cell map has been given by Jeffery (2007) and Moser, Kropff, and Moser (2008).

The Head Direction Cell Compass in Two Dimensions

Head direction cells typically show remarkably precise responsiveness to the orientation of the animal's head (not body), as shown in Figure 3.1d. They are found in a network of peri-hippocampal structures including lateral mammillary nuclei, anterior thalamus, dorsal pre-subiculum (also called post-subiculum), and retrosplenial cortex. Damage to any of these structures impairs navigation and also, for reasons that are not yet understood, impairs memory for life events in humans. Different cells prefer different head directions, and thus the cells can be thought of as a collection of tiny compasses, each one indicating a different direction. However, the population as a whole always acts coherently, which may allow a unique "sense of direction" to be possessed by the animal at any given moment.

A similar question arises with HD cells as for place cells: How does a HD cell manage to orient its firing so accurately? To compute one's orientation, it is necessary to have some kind of anchor that provides a reference direction, together with some means of updating computed heading following movement.

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

Visual landmarks seem to provide anchors for HD cells, because rotation of these reliably rotates the preferred firing directions (Taube et al. 1990b). These cues need to be located at some distance from the animal so that their angular location remains relatively stable as the animal moves around. Indeed, evidence shows that HD cells prefer distal landmarks (Zugaro et al. 2001) and that they perhaps do not use proximal ones at all, since place fields (an indirect measure of HD activity) do not rotate with rotation of intra-maze landmarks (Cressant et al. 1997). When an animal enters a novel environment, the newly encountered landmarks are, if possible, linked to the animal's internal sense of direction, which is maintained as it locomotes from familiar environment to new so that consistent orientations can be maintained across connected environments (Dudchenko and Zinyuk 2005). This linkage may not always occur; sometimes, the HD system as a whole may orient differently in different local environments (Taube and Burton 1995). It is possible that animals with magnetoreception or with a celestial compass might use this as a global signal to orient their HD system, but this has not yet been demonstrated.

To update their heading representation, evidence suggests that HD cells use multiple sources of angular self-motion information, both sensory and motor. On the sensory side, the vestibular apparatus detects angular movement with the semicircular canals, in the inner ear, which work by sensing the inertial flow of fluid induced when the head undergoes angular acceleration in any direction. The resulting signal is integrated to provide velocity and position information. In the horizontal plane, the HD cell signal breaks down following vestibular lesions (Stackman and Taube 1997), and thus this signal is clearly useful in updating the heading calculation of these cells. Other sensory sources of angular self-motion include optic and tactile flow, as well as proprioception. On the motor side, it is thought that a copy of motor commands is routed to the HD system to provide advance warning of an impending movement, enabling the system to respond almost immediately to the change in orientation. Evidence for this advance warning comes from the intriguing observation that the firing directions of HD cells in areas closer to the motor system, such as anterior thalamus, more accurately correspond to the animal's future head position than its current one; that is, they are slightly predictive, on the order of about 40 ms (Blair and Sharp 1995).

The Grid Cell Odometer in Two Dimensions

HD cells were discovered about twenty years after place cells, and it was another fifteen before the likely source of distance information to place cells was determined. Recording one synapse upstream of place cells, in entorhinal cortex, revealed spatially sensitive neurons whose firing is not singular, as with place cells, but occurs as repeating, evenly spaced circular regions spread in a hexagonal grid-like array across the surface of the environment (Hafting et al.

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

2005) (Figure 3.1e). These astonishing firing patterns strongly resemble graph paper (albeit hexagonal rather than square), and may indeed serve a similar function in the spatial system. Recently, evidence has emerged from an fMRI study for a grid cell network in humans as well (Doeller et al. 2010), suggesting generality of the processing mechanism across mammals.

Different grid cells lay down their fields at different positions, even if they are located adjacently in the brain; however, grids in a given animal in a given environment seem to have the same orientation, suggesting a coherent directional input, almost certainly arising from the HD system. The size of the spacing between the fields is characteristic of cells within a particular dorso-ventral level and is the same for all environments, implying the operation of an intrinsic distance-calculating process. Moving down the entorhinal cortex from dorsal to ventral, the scale of the grids becomes ever larger (Brun et al. 2008). Thus, the grid cell system has the capacity to represent environments of different scales. The stable orientation and distance properties make grid cells a likely candidate for the source of metric information to place cells. Consistent with this idea, place cells also increase their scale progressively from dorsal to ventral regions (Kjelstrup et al. 2008).

The question of exactly what kind of information the grid cells convey to place cells, and how, is as yet unresolved. Like place cells, grid cells also stretch and compress their fields as the environment is deformed (Barry et al. 2007). Therefore, they also appear to be anchored to the environment, perhaps by means of the more recently discovered "border cells" in entorhinal cortex (Solstad et al. 2008) and subiculum (Lever et al. 2009), whose firing lies along the boundaries of the environments. They also "remap" (change their firing patterns) when the contextual environment is changed (Fyhn et al. 2007; Alenda et al. 2010), again suggesting some consonance with the place field map. Thus, it seems likely that the activity of grid cells combines to form local, discrete firing fields in the place neurons. However, inactivation of the hippocampus (where the place cells are) causes the grid pattern to break down (Bonnevie et al. 2006), which makes a bidirectional effect seem likely. One possibility, proposed by O'Keefe and Burgess (2005), is that place cells feed their information back to the entorhinal cortex to anchor grid cells, and that grid cells in turn provide (by virtue of their evenly spaced fields) odometric information to place fields.

Place cells, HD cells, and grid cells have been well studied in flat environments, but the world, of course, is not actually flat. As noted earlier, all species have at least some vertical component to their travel, and those that can swim, climb, glide, or fly often spend much of their time traversing all three dimensions. An important issue, therefore, is how these cells operate when the animal moves into the third, vertical dimension. Based on findings from these neurons, the next section explores this question.

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

The Cognitive Map in Three Dimensions

Why should spatial representation in three dimensions be any different from navigation in two? There are several reasons. First, it is far more complex. In two dimensions, both heading (direction subject is facing) and bearing (direction subject is traveling) can be described with relatively few numbers, let's say, 360 at a resolution of one degree. In a volumetric space, this value, at the same resolution, increases to 130,000 (i.e., 360^2). A similar explosion in descriptive load occurs with the position parameter. Second, gravity means that navigation in the vertical dimension requires more energy; thus the vertical component of a journey makes a much bigger contribution to the cost of that journey. Third, to aid locomotion, the body plan of most animals is asymmetric in the vertical dimension (i.e., there is a "right way up" for the vast majority of living organisms), and this constrains how animals move and may therefore constrain how they map space. Finally, and perhaps most importantly, rotations in three dimensions have different properties from rotations in two in that they interact, which makes them order-dependent ("noncommutative"). This has implications for how a compass system could remain stable in three dimensions. As we will see, the brain may have solved this problem by keeping the vertical dimension separate and processing it differently.

Place Cells in Three Dimensions

Relatively few data exist concerning how place-coding neurons respond to excursions made into the vertical dimension. An early study of rat place cells, made on a surface that could be tilted (Knierim and McNaughton 2001), explored whether the fields would slowly shrink and disappear as the tilting surface moved through their (putatively) globular firing fields. This did not happen: although there was some slight responsiveness to slope, inasmuch as cells altered their firing patterns ("remapped") when the floor was tilted, there was no evidence of true 3D tuning.

Place fields, however, appear to show *some* sensitivity to height above ground, though less so than sensitivity to horizontal position. We made place cell recordings on both a helical track (akin to a spiral staircase) and a pegboard (like a climbing wall) (Figure 3.2; Haymann et al. 2011); in both kinds of apparatus, firing was modulated by how high the (horizontally oriented) animal was above the ground. Firing, however, extended over much greater vertical than horizontal distances, suggesting that the scale of the spatial map is greater in vertical space than horizontal space. Alternatively, it may be that place cells use nonmetric (and less-sensitive) encoding in the vertical dimension. It should be noted that in both kinds of apparatus, the animal's body maintained a horizontal posture while it was climbing. This may be important, as we will see below in the discussion on rotation.

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

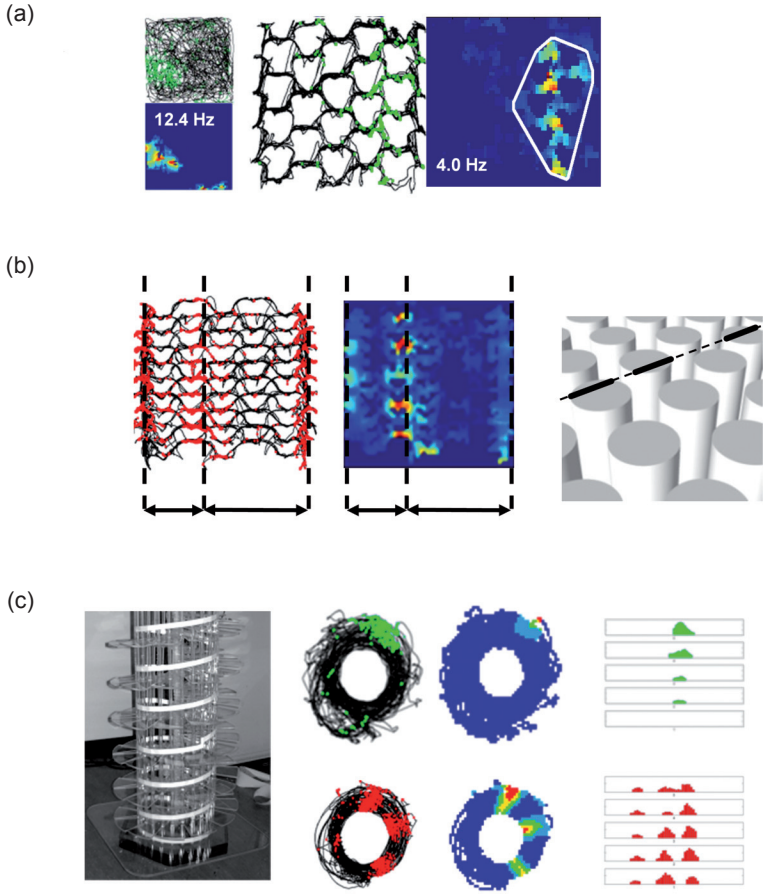

Figure 3.2 Responses of place and grid cells to travel in the vertical dimension, adapted from Hayman et al. (2011). (a) Place cell on the pegboard. Photo (left) shows a rat climbing on the pegs. Middle panel shows the raw data: the path of the rat is shown in gray and the cell spikes as small green spots. Right panel shows the same data as a contour plot (peak in red) with a white line to reveal the extent of the field. Note that the field shows vertical as well as horizontal delimitation. (b) A grid cell on the pegboard. Note that on the pegboard, the firing is aligned vertically as stripes. The middle panel shows the contour plot of the stripes, revealing more clearly the variable inter-stripe distance. The right panel shows a hypothesis about the origin of the stripes: namely, that they represent transections through an intrinsically columnar array of firing fields. (c) A place and a grid cell on a helical track (apparatus as shown in photo). The middle, circular plots show the track as seen from an overhead camera, showing the singular firing field of the place cell (top) and the multiple fields of the grid cell (bottom). When firing is decomposed into individual coils (right plots) and linearized to produce a firing rate histogram, it can be seen that the place cell's firing is vertically delimited whereas the grid cell seems to fire on all coils.

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

It remains to be seen whether place fields will turn out to be truly 3D (i.e., globular) in a volumetric environment where animals can move freely. There may also be ecological differences, in that animals that use surface locomotion (e.g., rats) may encode volumetric space differently than animals that can move freely in all dimensions (e.g., bats). Since ambulating bats evidently have both place cells (Ulanovsky and Moss 2007) and grid cells (Yartsev et al. 2010), telemetric recordings made as they fly will be particularly illuminating.

Head Direction Cells in Three Dimensions

Two issues must be addressed concerning directional orientation in 3D: How is heading initially computed by HD cells, and how is it updated following rotations? In 2D, heading can initially be established by carry-over of the internal direction sense (maintained by path integration) from the last known heading, or by static landmarks, or both. The same is probably true for 3D, though this has yet to be established. There is, however, an additional cue available in 3D, in the form of terrain slope, which can orient the place cell map (Jeffery et al. 2006) and therefore almost certainly the HD cells. The mechanism by which this acts is unclear. It could be a static cue, acting via gravity detectors, such as the otolith system in the vestibular inner ear together with visual and proprioceptive cues to surface inclination, or it could be a dynamic cue acting via proprioception and kinesthesis as the animal moves around. Experiments with moving versus stationary animals in light and in darkness are needed to address this question.

What about updating? The problem of updating heading following rotations in 3D is complicated. Rotations take place about an axis or, equally, within a plane, where the axis is orthogonal to the plane. This makes rotation a 2D process, unlike translation, with the consequence that rotation in a given plane has a component in each of the other two orthogonal planes. This is shown in Figure 3.3, where it can be seen that the translation of each point in the rotated object has a component in all three planes (X–Y, Y–Z, and X–Z).

Rotations in the three orthogonal planes are described in egocentric terms as "yaw," "pitch," and "roll." For an upright subject, as shown in Figure 3.4, yaw refers to rotations in the horizontal plane about a vertical axis, pitch to rotations in the sagittal plane about the transverse axis, and roll to rotations in the coronal plane about the antero-posterior axis. As there is no complete, convenient set of equivalent terms for rotations in an earth-centered frame of reference, the discussion below will use the egocentric terms together with additional information to orient the planes with respect to the environment.

Updating 3D heading signals has an added complexity because of these interactions between rotational planes; this means that a sequence of rotations in two orthogonal planes produces the same outcome as a single rotation in the remaining orthogonal plane. For example, the rat in Figure 3.5a began a sequence of turns by facing north and finished it by facing west, which amounts

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

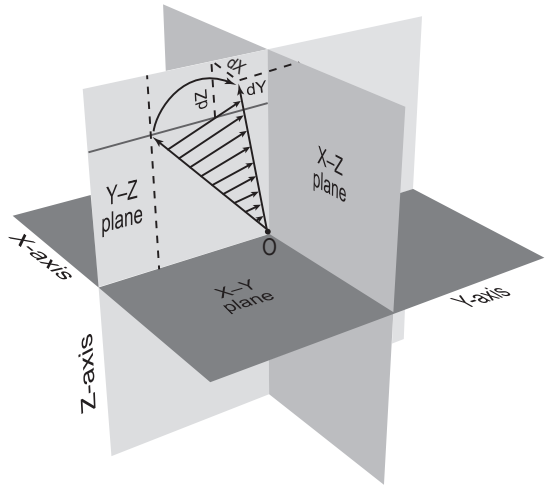

Figure 3.3 Rotations in 3D space. Each of the three spatial dimensions is referenced to one of three orthogonal planes, with each plane being defined by the two axes that bound it: X–Y, Y–Z and X–Z. The large arrow starts in the Y–Z plane and rotates (curved arrow) to its final position. The location of the arrow tip has thus moved in the X-axis (δX), Y-axis (δY) and Z-axis (δZ). If rotation is thought of as translation of a set of points (small arrows), where the translation distance for each point is proportional to its distance from the center of rotation (O), then it can be seen that the set of translations that comprise this rotation has a component in all three planes.

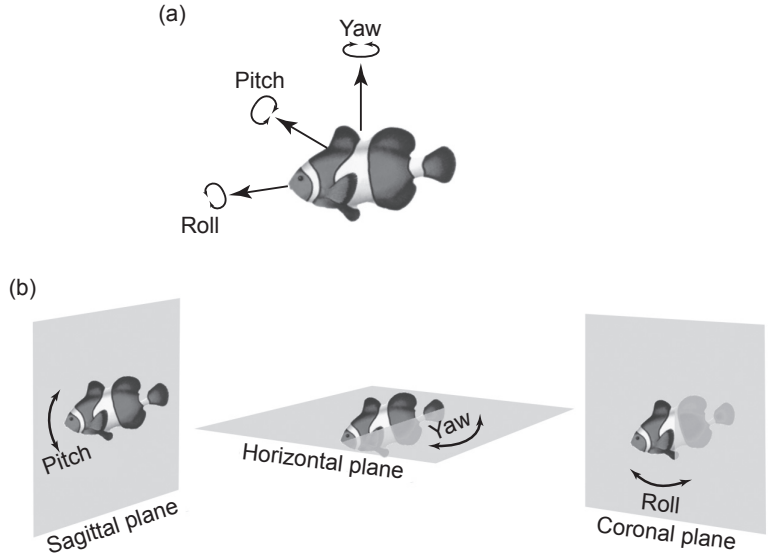

Figure 3.4 Rotational reference terms: (a) the three axes of rotation; (b) the three equivalent planes of rotation.

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

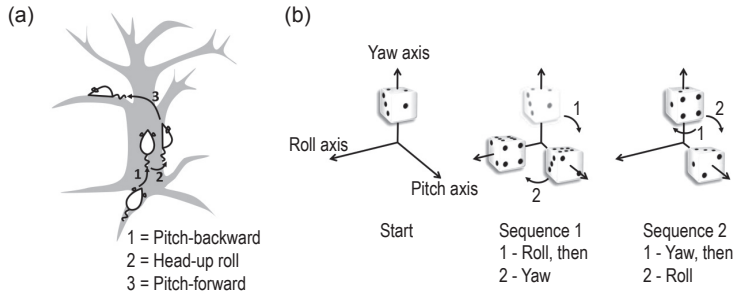

Figure 3.5 Interactions among 3D rotations. (a) Interaction of rotations in three dimensions as a rat climbs a tree, sidling around the trunk in the process. Although the animal has ended up on the branch with an orientation that represents a 90° counterclockwise yaw with respect to the starting heading, it never actually executed a yaw rotation. Thus, for HD cells to maintain a veridical representation of heading they need to adjust themselves during the vertical roll in step 2. This would not be the case, however, if the animal executed a horizontal roll (e.g., about the branch of a tree). Thus, the HD responses to roll need to be modulated by how the roll is oriented relative to vertical. (b) The noncommutativity of rotations in three dimensions. From a starting position (left), the die can undergo a roll followed by a yaw (middle), or a yaw followed by a roll (right). The different ordering of the two rotations results in a different final position and a different final orientation of the die.

to a counterclockwise yaw rotation, even though it did not execute a yaw at any stage of the process. Thus, the order of rotations matters: the end result is not the same for rotation *A* followed by *B* as it is for *B* followed by *A* (technically, it is said that rotations do not commute). This can be seen by rotating a polarized object, such as the die in Figure 3.5b, in 90° increments: first in the roll axis, then in the yaw axis, versus yaw followed by roll. Note that although the rotations in that example are described using allocentrically anchored axes, the same happens if the rotations are described in egocentric terms.

What does this mean for HD cells? It appears that HD cells do not provide a true 3D heading signal. Although, as we have seen, HD cells indicate heading in the horizontal plane and are sensitive to yaw rotations in this plane, they are not sensitive to pitch and roll rotations (Taube et al. 1990a). An experiment with rats climbing up a mesh wall on the inside of a cylinder found that HD cells maintained their firing as a rat moved from the horizontal to the vertical plane and thenceforth reacted on that plane as though the rotations were taking place on the horizontal plane (Stackman et al. 2000). This suggests that HD cells refer their firing to the animal's plane of locomotion and treat this as if it were horizontal. Firing on opposite sides of the cylinder was reversed: a cell that fired, for example, when the animal was facing directly up on one of the panels would fire when the animal was facing down on a panel on the opposite side of the apparatus, and both of these directions were also consistent with the firing on the arena floor. The cells thus did not treat the plane of locomotion in isolation; they also took into account the relationship of that plane to the arena

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

floor. The same geometric transformation that would map the arena floor onto each wall of the apparatus also described the transformation of the HD signal. Interestingly, the signal seemed to disappear completely when the animal was inverted (Calton and Taube 2005), and a behavioral experiment showed that animals also lost the ability to undertake map-based navigation while inverted (Valerio et al. 2010).

What does this mean for the animal's ability to orient in 3D? Clearly, HD cells do not suffice to allow this. There are cells elsewhere in the brain that are sensitive to pitch (Stackman and Taube 1998), but there is as yet no evidence that the signals from these are integrated with the HD signal to provide a true 3D compass heading. In any case, because the HD signal is not referenced to earth-horizontal, a complicated transformation would be needed to extract 3D heading from the mix of planar heading signals coming from the "horizontal" and vertical heading systems. Furthermore, the silencing of the HD signal in inverted locomotion suggests that the system is perhaps trying to avoid certain kinds of unstable state. In the next section, it is suggested that vertical and horizontal heading signals are kept separate to avoid the problem of noncommutativity.

Grid Cells in Three Dimensions

What about grid cells in 3D? The metric properties of grid cells offer a way to determine whether distance measurements in vertical space are encoded differently from horizontal space in the mammalian cognitive map. It might have been predicted a priori that grid cells would form grids in the vertical dimension as well as the horizontal, resulting in a lattice-like array of firing fields. In fact, recordings made on the helical track and the pegboard, described above for place cells, found that the cells showed their usual periodicity in the horizontal plane but, amazingly, did not seem to be modulated at all by height (or at least, very little). This produced a remarkable striped firing pattern when the animals were on the pegboard (Figure 3.6). Thus, it seems that grid cell odometry did not operate in the vertical dimension on these structures.

Two possible explanations may account for the failure of grid cells to account for vertical travel in this experiment. The first is that there is a distinction to be made between vertical odometry while an animal remains horizontal (as on the pegboard and the helix) and odometry when the animal's plane of locomotion (i.e., the plane of its feet) is also vertical (e.g., when climbing a mesh), and it may be that the limitation to vertical odometry applies to the former but not to the latter. In other words, perhaps the odometer cannot determine distances in the direction orthogonal to the plane of the animal's body. Second, it may be that the odometer does not measure distances in the direction of gravity, but only distances orthogonal to gravity (i.e., in the earth-horizontal plane).

If grid cells are not measuring vertical distances traversed, what *are* they measuring? Do they measure distance traveled in the earth-horizontal plane

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

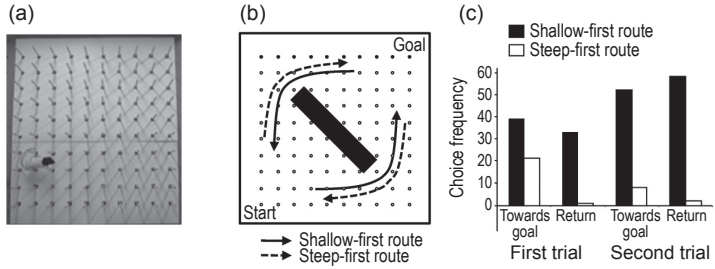

Figure 3.6 Anisotropic route choice in rats, from an experiment reported in Jovalekic et al. (2011). Rats were trained on the pegboard (a) to shuttle between a start and goal corner. When a barrier was introduced, rats had the choice of taking either a route with a steep segment first followed by a shallow segment, or a shallow segment first followed by a steep segment (b). Rats strongly preferred the shallow-first route (c) on both test trials.

alone, or do they measure surface distance traversed, which would be longer on, say, hilly terrain? The limited available experimental data on 3D path integration predict both possibilities; optic flow and locomotor cues (e.g., footstep counting) convey information about surface travel, whereas data from HD cells suggest that the frame of reference, as discussed above, may be the plane of locomotion rather than earth-horizontal (Stackman et al. 2000). On the other hand, the vestibular apparatus encodes movements relative to the inertial frame of reference of the earth's gravitational field, and data from insects suggests extraction of the earth-horizontal component of travel (Wohlgemuth et al. 2001). Thus, the question has not yet been resolved.

Is the Cognitive Map Anisotropic?

Based on the evidence from place-coding neurons reviewed above, it is hypothesized here that the brain's representation of allocentric space—the cognitive map—is anisotropic, encoding vertical space differently from horizontal space. Specifically, it is proposed that horizontal space is encoded metrically and vertical space nonmetrically. This hypothesis derives from observations that (a) HD cells do not show sensitivity to pitch rotations, and (b) place and grid cells appear to encode vertical space differently, inasmuch as both cell types show elongated firing fields and grid cells do not show periodicity in the vertical dimension in the same way that they do in the horizontal. Since grid cell periodicity is thought to be the means by which place cells receive distance and direction information, the absence of vertical periodicity in grid cells suggests that the place cell map may not itself be metric in the vertical dimension. Thus, the tentative conclusion is that the cognitive map measures distances horizontally but not vertically. Note that "horizontal" here assumes that the animal is oriented normally (i.e., horizontally), moving around with

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

its feet on a plane parallel to the earth's surface. The situation for climbing a vertical wall may, as discussed above, be different, and this issue needs further investigation.

The anisotropy hypothesis is only tentative at present, because place cells may receive metric inputs other than grid cells, which have not been discovered yet. However, map anisotropy is supported by behavioral studies in both animals and humans. Studies of rats in a 3D lattice maze found that they tended to solve the vertical and horizontal components of the task independently (Grobéty and Schenk 1992), and we have obtained similar findings both in a lattice maze and on the pegboard (Jovalekic et al. 2011). Interestingly, in our study we found that the horizontal components of the tasks tended to be solved first: when offered a choice of flat-steep versus steep-flat routes to the same place, rats invariably chose the flat-steep option, irrespective of whether they were going upward or downward (Figure 3.6), suggesting a preference for executing the less costly part of the journey first. A conceptually similar study in humans found that subjects solved the horizontal part of a journey first in a multilevel building (Hölscher et al. 2006). These effects might not be confined to surface-dwelling animals or to mammals. Holbrook and Burt de Perera (2009) found that fish would separate the vertical from horizontal components of a spatial choice. In their experiment, fish were trained on a Y-maze to choose between a left up-slanted tube or a right down-slanted one, and then tested with the conditions reversed (left-down and right-up). The fish tended, surprisingly, to choose the tube that matched the training condition in its slope characteristics regardless of whether it lay to the left or the right. This suggests not only an ability to separate the vertical from horizontal dimensions, but also a preference for vertical information as being perhaps more salient.

Studies of humans orienting in 3D space find that subjects tend to prefer to reduce the degrees of freedom in their representation. For example, Vidal et al. (2004) found that subjects in a virtual 3D maze performed better if they remained upright and had therefore aligned their egocentric and allocentric frames of reference in one of the dimensions. Similarly, astronauts in a weightless environment tend to construct a vertical reference using visual rather than vestibular (i.e., gravitational) cues and remain oriented relative to this (Tafforin and Campin 1994).

What could be the advantages and disadvantages of an anisotropic map in which horizontal and vertical components are separated? At first glance, this seems like a maladaptive strategy. Although a 2D orientation system is commutative in its operation, by ignoring the third dimension the system is vulnerable to mistakes, in situations like the one shown in Figure 3.5a. Another disadvantage is that 3D vectors cannot be computed with a map that is only metric in the horizontal plane. This can be seen by considering a contour map, in which horizontal information is represented in Cartesian (X–Y) coordinates while vertical information is represented by numbered lines. Although it is possible to locate one's position in 3D on such a map, it is not easy to perform

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.

navigational computations such as vector addition, needed for homing, or detour planning.

On the other hand, there are several advantages to allocation of horizontal and vertical components of the map to separate processors. To begin, a planar rather than volumetric map imposes a much smaller representational load (e.g., 360 vs. 130,000 "bins" in the directional example given earlier). Second, the separation of horizontal from vertical means that each component of the map only has, at most, a 2D problem to solve. The horizontal, metric component can compute rotations using a commutative process (e.g., simple integration of successive angular displacements) without needing to process the order in which these occur. The vertical processor only needs some kind of information about height: a one-dimensional property. Unless this needs to be integrated with the horizontal processor for 3D vector calculation, a fairly coarse, non-metric representation of height would suffice for most purposes, at least for surface-traveling animals such as rats and mice (and humans). This is because on a surface, the vertical dimension is yoked to the horizontal: for any given X–Y position, there is only one Z position. Therfore, it may be only necessary to have a coarse representation of height to allow heuristic route calculations to occur (e.g., if route *A* has more hills than route *B* for the same horizontal distance, then select route *A*). Whether or not surface-traveling animals perform precise computations of height traveled, or only use heuristic approximations, is a matter for future research.

Conclusions

This chapter hypothesizes that the cognitive map is anisotropic; that is, the brain has evolved to process horizontal and vertical dimensions separately to reduce representational and computational load. Testing this hypothesis will require studies in both behavioral and neural domains. If confirmed, it will have implications not only for the academic study of cognitive mapping, memory, and executive functions, but also technological design in three dimensions, including architecture, spacecraft design, and virtual reality.

Acknowledgments

This work was supported by grants from the Wellcome Trust (GR083540AIA) and the European Commission Framework 7 ("Spacebrain"). The author would like to thank Theresa Burt de Perera and Eleonore Haussner for helpful comments on a draft of the manuscript.

From "Animal Thinking: Contemporary Issues in Comparative Cognition," edited by R. Menzel and J. Fischer. Strüngmann Forum Report, vol. 8, J. Lupp, series ed. Cambridge, MA: MIT Press. ISBN 978-0-262-01663-6.