



Spatial Cells in the Hippocampal Formation

Nobel Lecture, 7 December 2014

by John O'Keefe

University College London, United Kingdom.

INTRODUCTION

When I first started working on the hippocampus in 1970, the most important clue we had to its function had come from the work of Brenda Milner (Figure 1B) on the patient Henry Molaison (Figure 1A) (Scoville and Milner, 1957). Henry had undergone extensive removal of both medial temporal lobes including the hippocampus as well as surrounding structures for the relief of severe temporal lobe epilepsy (Figure 1C, 2A). While the epileptic seizures were ameliorated, the most dramatic result of the operation was a profound anterograde and retrograde amnesia. He not only failed to lay down new memories of his day-to-day experiences but was unable to recall ones previously stored. He suffered a profound episodic memory deficit. In the words of Sue Corkin (Figure 1D), one of Brenda's students who studied him extensively over the years

He . . . cannot recall anything that relied on personal experience, such as a specific Christmas gift his father had given him. He retained only the gist of personally experienced events, plain facts but no recollection of specific episodes (Corkin, 2013, p. 219).

In the late 1960s, I had been recording somatosensory cells in the thalamus of the rat using our newly developed techniques for recording from single units in freely moving animals and on one occasion had inadvertently positioned electrodes more laterally in the hippocampus. There I found a cell which had a beautiful phase relationship to the 7–10 Hz local field potential theta activity

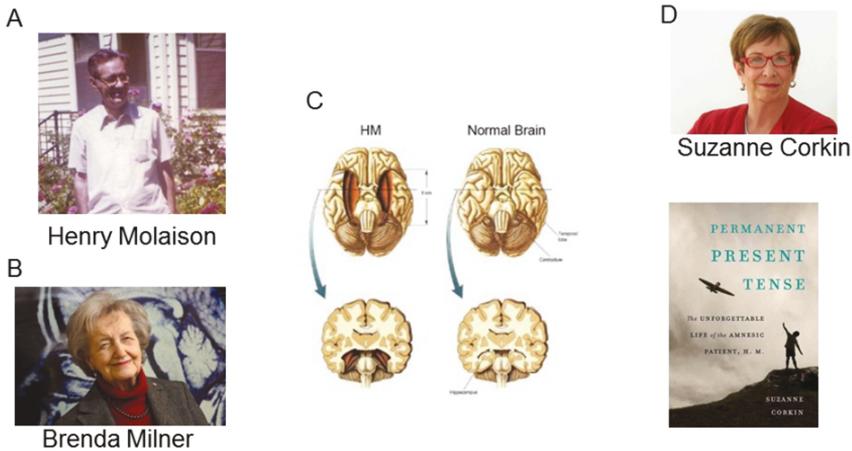


FIGURE 1. Henry Molaison (HM) (A) lost his episodic memory following medial temporal lobectomy (C). His amnesia was initially reported and studied by Brenda Milner (B) and more recently by Suzanne Corkin (D) as reported in her recent book “Permanent Present Tense.”

and which clearly was activated by some higher-order aspect of the animal’s movement such as speed of movement of the head. This behavioural correlate intrigued me since, on the face of it, it seemed very far removed from the type of cell you might expect to find in a memory system. I immediately decided to abandon the somatosensory system and move to the study of the hippocampus in an attempt to see what memories looked like at the single cell level. The hippocampus in the rat has very much the same anatomical structure (Figure 2B) as in the human and I reasoned that we might be able to gain considerable insight into how memories are stored and retrieved. Around this time I was joined in the lab by Jonathan Dostrovsky, an MSc student, and we decided to record from electrodes in the CA1 field (Figure 2B) as the animal performed simple memory tasks and otherwise went about its daily business. I have to say that at this stage we were very catholic in our approach and expectations and were prepared to see that the cells fire to all types of situations and all types of memories. What we found instead was unexpected and very exciting.

Over the course of several months of watching the animals behave while simultaneously listening to and monitoring hippocampal cell activity it became clear that there were two types of cells, the first similar to the one I had originally seen which had as its major correlate some non-specific higher-order aspect of movements, and the second a much more silent type which only sprang into activity at irregular intervals and whose correlate was much more difficult to

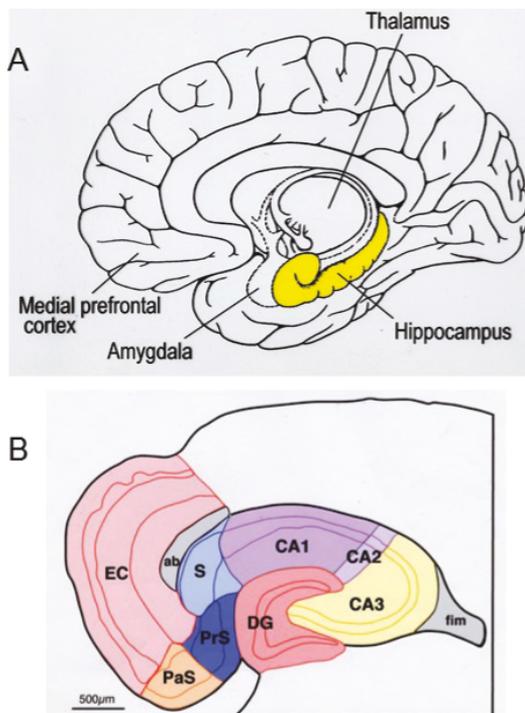


FIGURE 2. The human hippocampus is located in the medial temporal lobes (A). The rodent hippocampal formation (B) is composed of the hippocampus proper (CA fields and the dentate gyrus DG), the subicular region (subiculum S, pre-subiculum PrS, and para-subiculum PaS) and the entorhinal cortex EC.

identify. Looking back at the notes from this period it is clear that there were hints that the animal's location was important but it was only on a particular day when we were recording from a very clear well isolated cell with a clear correlate that it dawned on me that these cells weren't particularly interested in what the animal was doing or why it was doing it but rather they were interested in where it was in the environment at the time (Figure 3). The cells were coding for the animal's location! In the short three page paper (O'Keefe and Dostrovsky, 1971) that we wrote to summarise our initial findings, we said

These findings suggest that the hippocampus provides the rest of the brain with a spatial reference map. Deprived of this map . . . it could not learn to go from where it happened to be in the environment to a particular place independently of any particular route (as in Tolman's experiments) . . . (pp. 174–5)

EC Tolman had been one of the doyens of American behavioural psychology who believed that animals found their way around environments by creating internal representations, which were more complicated and interesting than the simple associations between stimuli and responses beloved of the behaviourists of the Hullian persuasion. Tolman wrote:

We believe that in the course of learning, something like a field map of the environment gets established in the rat's brain . . . The stimuli . . . are usually worked over . . . into a tentative, cognitive-like map of the environment. And it is this tentative map, indicating routes and paths and environmental relationships, which finally determines what responses, if any, the animal will finally release. (Tolman, 1948, p. 192).

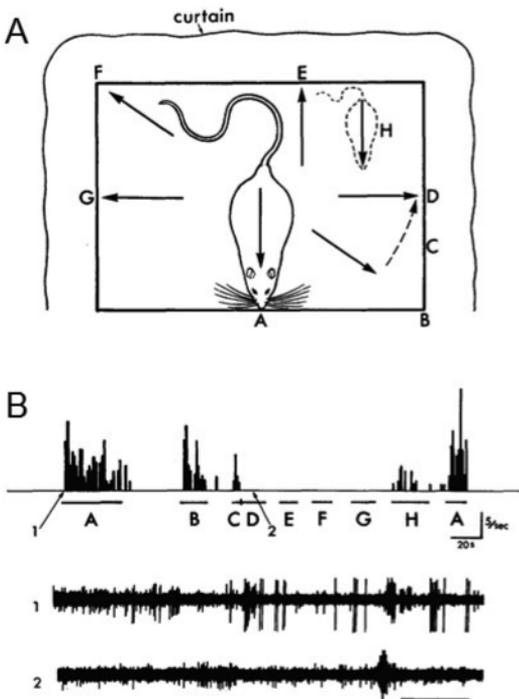


FIGURE 3. The first report of a place cell in the hippocampus of the freely moving rat. The animal went to different locations and faced in different directions on a rectangular platform surrounded on three sides by curtains (A). The cell fired when the animal was in locations A and B but not elsewhere. After O'Keefe and Dostrovsky (1971).

He never specified in any detail what this map would look like or consist of, and it is highly unlikely that he thought it would be localised in one brain system rather than being distributed throughout the whole brain. He also expected it to be involved in much more than spatial and navigational functions. This together with the absence of any techniques for identifying the cognitive map and its components led the idea to fall into disuse and the argument to be won by the Hullian behaviourists. But here Dostrovsky and I were recording from cells that looked like one important component of such a map, representing as they did different places in a familiar environment. Possessing such a map would surely furnish an animal with a very handy cognitive device enabling it to move around environments in a creative, flexible way rather than being reduced to the simple rigid stimulus-response routes envisaged by the behaviourists. This thinking led naturally to the prediction that damage to the hippocampus in rodents would result in a specific set of deficits in spatial location and spatial navigation as embodied in the quote from the 1971 paper. It also dawned on me that maps were rather sophisticated devices and needed more spatial information than that embodied by just a set of place representations. For one thing, it seemed likely that a place could be identified in at least two independent ways. In addition to a computation based on the sensory stimuli impinging on the animal when it found itself in that place it might be possible to compute that the animal had arrived at the same location on the basis of an updating of its position using information from its own movements, a technique not dissimilar to that used in the navigation systems of aeroplanes. In the subsequent paper in which I described the properties of place and misplace cells more extensively, I speculated that

Each place cell receives two different inputs, one conveying information about a large number of environmental stimuli or events, and the other from a navigational system which calculates where an animal is in an environment independently of the stimuli impinging on it at that moment . . . When an animal had located itself in an environment (using environmental stimuli) the hippocampus could calculate subsequent positions in that environment on the basis of how far and in what direction the animal had moved in the interim . . . (O'Keefe, 1976)

And in a summary of the contents of a book "Hippocampus as a Cognitive Map"

A small number of stimuli (two or three) occurring with a unique spatial configuration when an animal is in a particular part of an environment are sufficient to identify a place in the map. At any given point in an environment there are usually a large number of such sets of stimuli and therefore the identification of a place in an environment does not depend on any particular cue or group of cues. The distance and direction vectors which connect the places in the map of an environment are derived from the animal's movements in that environment. (O'Keefe and Nadel, 1979, p. 489).

An idea of what we had in mind in these quotes is shown in Figure 4 taken from a subsequent paper (O'Keefe, 1996) in which the place representations A, B and C are connected by vectors representing the distance between them in a certain direction. This type of representation would be useful in two ways: Firstly, if an animal found itself at location A in a familiar environment and wanted to go to location B (perhaps because it was hungry and had a representation that there was food there) then the system would generate the vector AB and send it to the motor system causing the animal to travel in a specific direction for a specific distance. The flexibility in the system arises from the fact that under circumstances in which the direct route AB is blocked, the system can generate a detour using the fact that the sum of vectors AC and CB equal AB. A second use of the system is in predictive mode. If the animal finds itself travelling along vector AB from location A, the mapping system will predict that it will end up at location B and furthermore what it should experience there. This ability of the internal navigation system to generate predictions about what will be experienced at particular locations forms the basis of learning in the cognitive map. If there is a mismatch between what the system predicts at B and what the animal experiences at B, a mismatch signal is generated which initiates

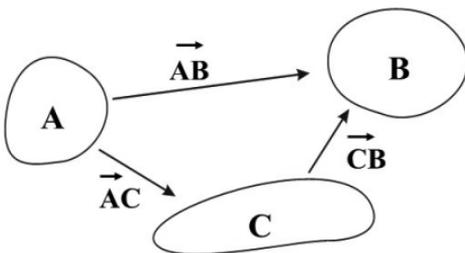


FIGURE 4. Schematic representation of a cognitive map of three places (A, B, C) and the vectors connecting them (AB, AC, CB). After O'Keefe (1996).

exploration behaviour designed to initially build maps of an unknown environment or to update existing maps where they are no longer veridical, perhaps due to small changes in the actual environment. In the 1976 paper, I reported the existence of a subset of place cells which I called misplace cells and which fired maximally when the animal went to a familiar place and found an unexpected novel object there or failed to find an expected familiar object. The ability of the mapping system to identify locations on the basis of subsets of environmental cues, to predict the sensory cues which would be found in particular locations, and flexible navigation involving the generation of both shortcuts and “umwegs” (detours) provide a powerful cognitive mechanism for learning about and representing an environment, and flexibly navigating through it.

Around this time, I was joined by Lynn Nadel in the Anatomy Department at UCL (Figure 5 right) and we decided to write a review paper showing how hippocampal anatomy and physiology could underpin the predicated mapping system and how the loss of such a map could explain the behavioural deficits following hippocampal damage. The review paper grew into a sizeable book and after a long gestation period was eventually published in 1978 by Oxford University press under the title “The Hippocampus as a Cognitive Map” (Figure 5 left). In it we covered the extant literature on the hippocampus, but in addition we tried to set our ideas in their historical context by covering the philosophical and psychological literature on space and spatial representation. We were deeply impressed by how important to many different behaviours such a spatial system might be and waxed lyrically about the role of space as shown in the poetic language with which we introduced the book.

THE
HIPPOCAMPUS
AS A COGNITIVE MAP

JOHN O'KEEFE
AND
LYNN NADEL



CLARENDON PRESS · OXFORD



Lynn Nadel

FIGURE 5. Lynn Nadel and John O'Keefe kibitzing during the writing of the Hippocampus as a Cognitive Map.

SPACE

plays a role in all our behaviour.

We live in it, move through it, explore it, defend it.

We find it easy enough to point to bits of it:

the room,
the mantle of the heavens,
the gap between two fingers,
the place left behind when the piano
finally gets moved.

(O'Keefe and Nadel, 1978, p. 5).

Theories make predictions which go beyond the extant data. We predicted the existence of hippocampal signals coding for direction, distance and speed of movement and showed how the known effects of hippocampal lesions could be explained by impaired place learning, navigation, and exploration. Having reviewed the lesion literature, we also suggested that most existing behavioural tasks used to test hippocampal function were inadequate since they did not isolate the spatial components of the task sufficiently well to demonstrate that they were the crucial components. Further they did not adequately rule out alternative strategies based on remaining brain systems so that e.g., the animals couldn't learn using stimulus/response route-based strategies. The book had a limited production run and was not readily available for many years. To rectify this situation, some years ago I scanned the original book and placed the scanned version on the Internet (www.cognitivemap.net) to make it more accessible.

Well, if you make predictions and are lucky enough to have colleagues who take them seriously, they will set out to test them. One of the first persons to test a prediction of the theory was Richard Morris (Figure 6A). Inspired by the cognitive map idea, Richard who was then at St Andrews University devised a simple but powerful spatial navigation task which required the animal to approach a location defined by distal environmental cues when started from any one of several different start positions thus forcing it to approach the goal from different directions on different trials. The "Morris water maze," as it came to be known, consisted of a tub of opaque water within which was placed a hidden sunken escape platform effectively eliminating local cues (Figure 6B). Richard did all the important groundwork showing how normal animals performed on the maze and controlling for the influence of all but spatial information (Morris,

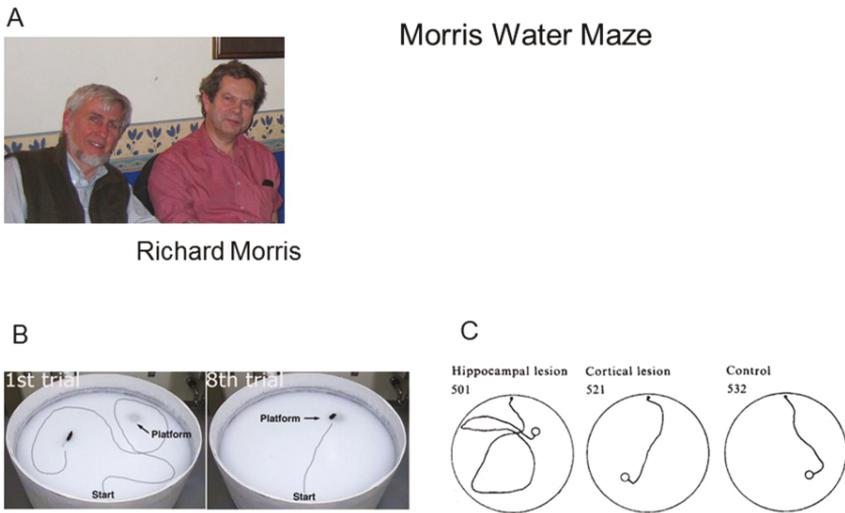


FIGURE 6. The Morris water maze was developed by Richard Morris (A) in order to test predictions of the cognitive map theory. The animal is placed in a tank filled with opaque water containing a hidden platform which it must locate on the basis of extramaze room cues (B). Control animals and those with neocortical lesions readily learn to do this but animals with damage to the hippocampus find it much more difficult (C).

1981). Subsequently Richard Morris, Nick Rawlins from Oxford and ourselves tested animals with hippocampal lesions on the task and showed that, unlike control animals and those with lesions of the overlying neocortex, these animals were severely deficient in learning to find the hidden platform (Figure 6C). The water maze has proved an invaluable test of hippocampal function and is now one of the most widely used behavioural testing apparatuses in neuroscience.

In the late 70s I was lucky enough to attract Bruce McNaughton and Carol Barnes to spend a year as postdocs in my laboratory (Figure 7 right). They were already experts in intracellular recording, long-term potentiation and behavioural studies. We had just introduced a minicomputer and an overhead camera head tracking system to the lab and Bruce, Carol and I carried out the first quantitative measurement of place fields albeit on the one-dimensional arms of an 8-arm maze (McNaughton et al., 1983). We also showed quantitatively for the first time that the firing rate of place cells was dependent on the animal's speed (Figure 7A). It was an important prediction of the cognitive map theory that the animal's speed would be represented in the hippocampus because this speed signal was needed to translate movements through the environment at different speeds into shifts of the same displacement within the map representation. Subsequently we were able to find a small number of pure speed cells in

Place Cell Firing Rate Modulated with Speed

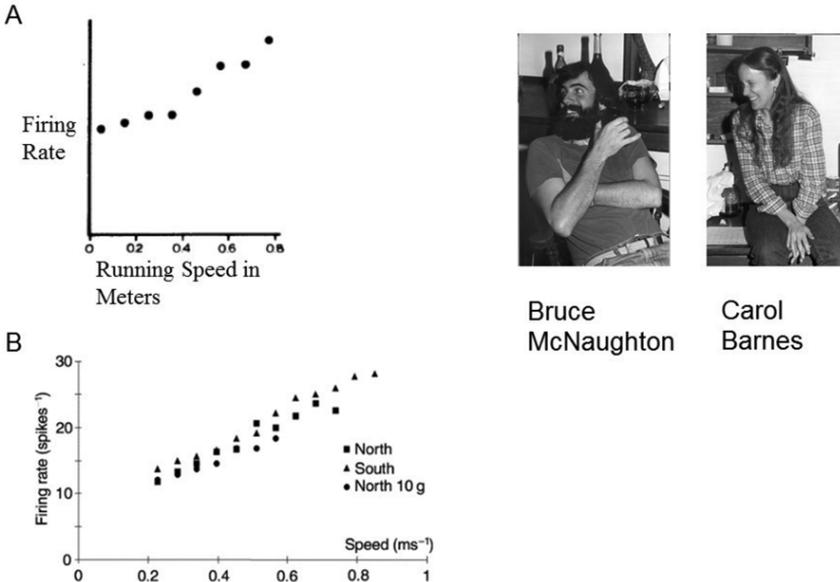


FIGURE 7. Place cell firing rate is modulated by the animal's speed of running. Bruce McNaughton and Carol Barnes (upper right) showed a good correlation between the two measures (A). Subsequently a small number of pure speed cells have been found in the hippocampus itself (B). Firing in these cells is not dependent on direction of movement or on the amount of exertion or force required to carry out the movement. After McNaughton et al. (1983) and O'Keefe et al. (1998).

the hippocampus which fired irrespective of the animal's location or direction of movement and which did not change markedly if the animal was pulling a 10 gram weight, ruling out a role for muscular effort (O'Keefe et al., 1998) (Figure 7B).

WHAT ABOUT THE OTHER CELL TYPES REQUIRED BY THE THEORY?

We now know that all the cell types required by the cognitive map theory exist in the greater hippocampal formation (Figure 8): place cells in the hippocampus proper which tell the animal where it is in a familiar environment (A), head direction cells in the presubiculum and entorhinal cortex which tell the animal which direction it is pointing in (D), boundary cells in the subiculum and entorhinal cortex which tell how close it is to a boundary in a particular direction (B), and last but not least, grid cells in the entorhinal cortex which appeared to

Spatial cells in the hippocampal formation

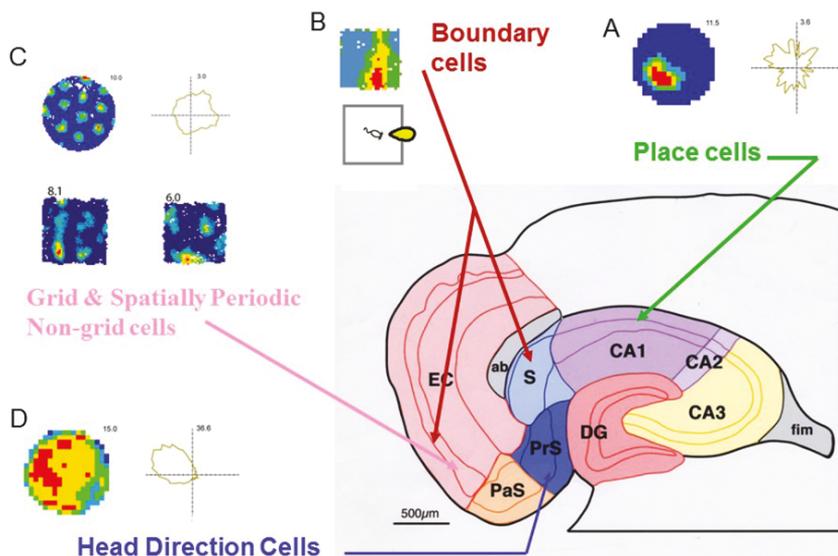


FIGURE 8. Spatial cells in the hippocampal formation include (A) place cells of the hippocampus proper which fire in localised parts of the environment (firing rate map, left) and in open fields irrespective of the animal's heading direction (polar plot, right), (B) boundary cells of the subiculum which fire when the animal is located a distance from a boundary in an allocentric direction, (C) grid and other spatially periodic cells of the medial entorhinal cortex which fire in multiple locations across an environment in a symmetrical hexagonal pattern (grids) or in other symmetrical but non-hexagonal patterns (spatially periodic non-grid cells), and (D) head direction cells of the pre-subiculum which fire when the animal's head is pointing in an allocentric direction (polar plot, right) irrespective of location in the environment (firing rate map, left). Boundary cells are also found in the entorhinal cortex where they are called border cells and head direction cells are also found in the entorhinal cortex. After Lever et al. (2009) & Krupic et al. (2012).

provide the metric for the map at least in some environments (C). In addition to the grid cells there are several other cell types in the entorhinal cortex including spatially periodic non-grid cells which fire in repeated regular patterns without the hexagonal symmetry of the grid cells and which may carry out some of the same functions as the grids (Krupic et al., 2012). In this review I can't possibly attempt to summarise all we know about these different cell types but will point out a few important aspects of each in the next few sections.

Place cells are typically recorded in boxes with square or circular walls (Figure 9A) where the animal is free to move in all directions except when it is close to the walls. In these environments, the place cells typically have one localised

Place cells and cognitive maps

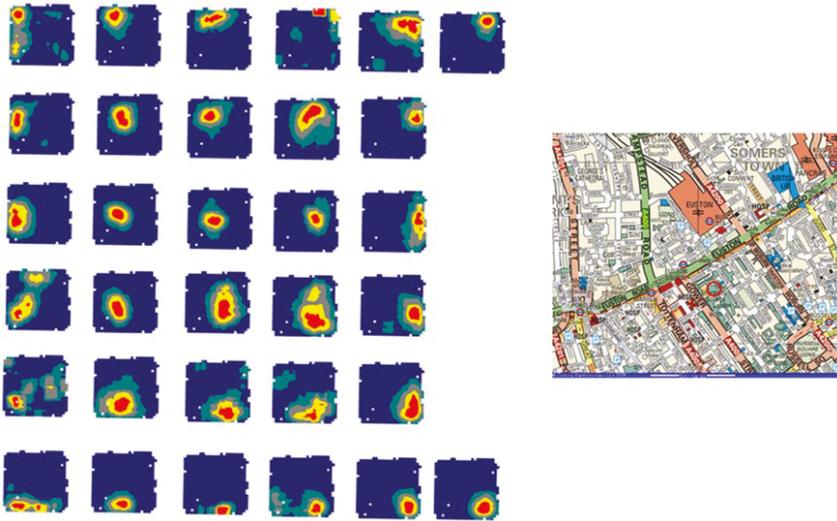


FIGURE 9. Different place cells have different firing fields and 32 simultaneously-recorded place fields taken together cover the surface of an environment. Firing rates are arranged to correspond to field locations in the environment and not to anatomical location in the hippocampus.

firing field within which it is happy to fire independently of the animal's direction of facing or movement. In environments in which the animal's movements are constrained to move in one or two directions (such as linear tracks, see Figure 12), the place cells typically have a preferred direction and fire much less or not at all in the opposite direction (McNaughton et al., 1983; O'Keefe and Recce, 1993). When a group of cells are recorded at the same time, each has its own firing field and, taken together, the firing fields of even a small group of cells tend to cover the entire surface of the environment (Figure 9). Place cells recorded in two familiar discriminable environments such as a square and a circle clearly discriminate between the two environments (Figure 10). Some fire in the square but not the circle (top row), others fire in the reverse pattern (second row), and still others fire in both boxes but not in the same locations (third row), and finally (bottom row) a small group fire in the same location in both boxes probably because the boxes were located in the same place in the laboratory and the cells were firing in the laboratory frame. Place cells recorded in three-dimensional environments have three-dimensional fields. The best demonstration of this is in bats, where the fields tend to be spherical and also cover the volume

Place cells differentiate between 2 environments

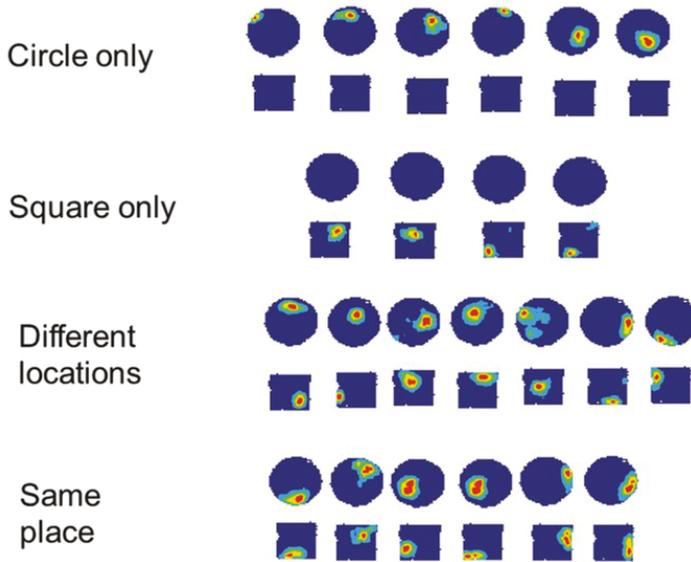


FIGURE 10. Place cells differentiate between 2 familiar environments, a circle and a square, either by firing only in the circle (top) or the square (2nd row), firing in both but in different locations (3+ row). A small number of cells fire in the same place in both boxes (bottom) probably because the boxes were set in the same location relative to the experimental room and the cells were responding to the absolute location in the room.

of space in the same way as two-dimensional place cells in the rat (Figure 11) (Yartsev and Ulanovsky, 2013).

Temporal coding in place cells allows place cells to identify locations in an environment with a much finer precision than if only firing rates are used (Figure 12). To know when an action potential has occurred, you need some sort of clock against which to measure time. The way the place cells do this is by using the population activity of the hippocampus as represented by the sinusoidal local field potential (LFP) called theta. The local field potential recorded from an electrode in the hippocampus shows a striking rhythmical oscillation called the theta rhythm which typically varies in frequency from 6 to 10 Hz (Figure 12A). When this oscillatory activity was first described, many thought it was an artefact since it didn't really seem likely that the brain was creating sine waves. In fact many years of experimentation have shown that theta is not only real but also a very important aspect of hippocampal function. Furthermore there is a nice behavioural correlate of theta as originally shown by Case Vanderwolf

3-Dimensional Place Fields

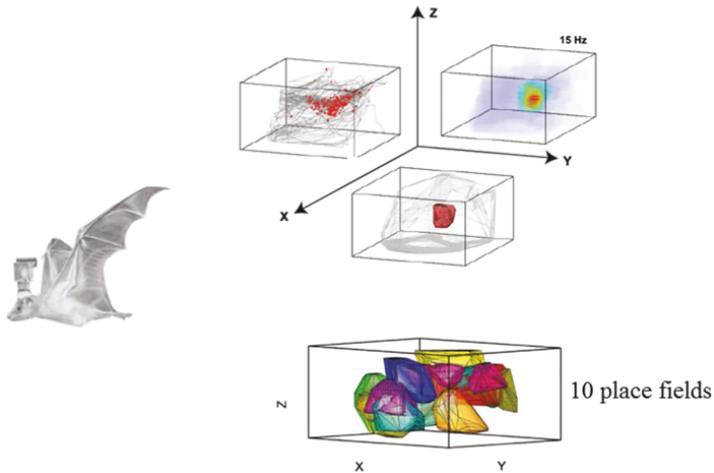


FIGURE 11. Three-dimensional place fields are found in the hippocampus of flying bats. The spherical fields of 10 place cells fill up the volume in which the animal is navigating. After Yartsev et al. (2013).

(Vanderwolf, 1969). It occurs primarily when the rat runs around an environment engaging in behaviours which translate its location relative to the environment (Figure 12A). In contrast when the animal sits still during drinking or eating, the rhythm typically disappears to be replaced by a less organised pattern with a broader range of frequencies. Over the years, Gyuri Buzsáki (Figure 12 upper left), who has studied hippocampal theta activity more than anyone else, would occasionally ask me about the relationship between the firing pattern of place cells and theta waves. We both knew from the earliest recording experiments that the interneuronal theta cells tended to have a fixed phase relationship to a particular part of the theta wave. In contrast I had not been able to identify any such fixed phase relationships between place cell firing and theta in earlier informal recordings and told him so more than once. Thinking about this around 1990, it seemed improbable that there would be no relationship whatsoever, so I decided to look more closely at the relationship on an individual run-by-run basis. I used a dataset that had been collected by Andrew Speakman and myself in a memory experiment carried out on a 4-armed plus-shaped maze with controlled environmental cues. The study was originally designed to monitor the firing of place cells when the animal was allowed to locate itself relative to the environmental cues but then asked to navigate to a goal after these cues

Temporal Coding of Location



Gyuri
Buzsaki



Mike Recce

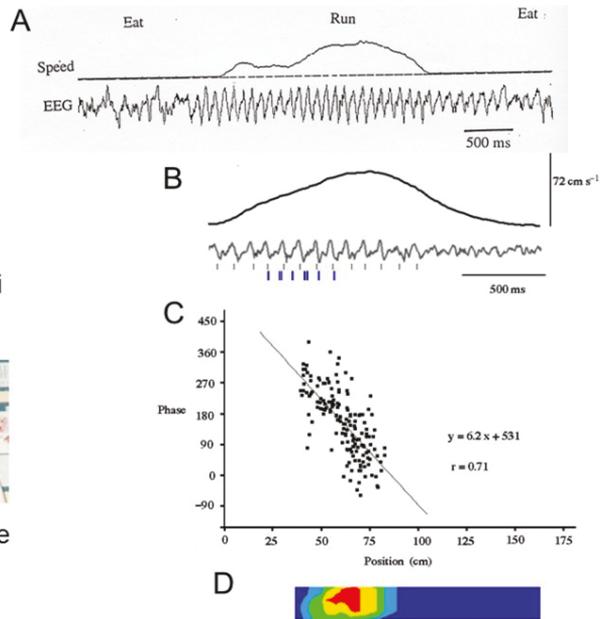


FIGURE 12. The timing of place cell firing as well as rate is used to code spatial information. Hippocampal local field potential shows 8–10 Hz theta oscillations as the animal runs in a linear track but a less organised irregular pattern as the animal sits at the end of the track eating (A). When the animal runs through the field of a place cell (D), the place cell emits a series of action potentials (B, blue) whose frequency is slightly higher than the simultaneously recorded LFP theta (B, black ticks). The result is that the spikes precess relative to the theta waves occurring at earlier phases each successive burst. The phase of place cell firing is highly correlated with the position of the animal (C). After O'Keefe and Recce (1993).

had been removed (O'Keefe and Speakman, 1987). I concentrated on the relationship between a well-isolated familiar place cell and the ongoing theta activity. After several days of pondering over what seemed at first to be a random relationship in which the spikes could fire at any phase of the theta wave, I realised that the relationship was not random but systematic: as the animal ran through the field each successive burst of spikes moved earlier and earlier relative to the theta wave. The spikes were precessing in a systematic way relative to the theta wave (Figure 12B). Notice how the blue spike moves earlier and earlier relative to the black ticks marking the + to -0° crossing of the theta wave as the animal runs through the place field. It seemed possible that this temporal code was carrying information. To study this more quantitatively, Michael Recce (Figure 12 lower left) and I recorded the place cells and theta waves as rats ran back and

forth between the two ends of a linear track for a food reward. A typical instance of what we saw is shown in Figure 12C. When we plotted the phase of firing of the cell against the animal's location on the linear track we found a striking correlation suggesting that position is coded not only by the rate of firing of the cell but also by the exact timing of spikes relative to this wave. To this day, this phase coding for location is one of the best examples of a temporal code in the nervous system. More generally, temporal coding may be a very important property of cells in the cortex. In subsequent work with Neil Burgess and John Huxter (Huxter et al., 2003), we showed that the rate and temporal codes were relatively independent, opening up the possibility that the two variables might code for different aspects of the animal's behaviour or location. This might be one solution to the binding problem: how does the nervous system know that two different variables go together? Perhaps they are represented by different aspects of the same train of action potentials in a single cell.

Boundary cells may be part of the mechanism for identifying places in the environment on the basis of sensory information impinging on the animal at particular locations. The first clue to the existence of these cells came from an experiment that Neil Burgess (Figure 13 right) and I carried out on place cells in the mid-1990s. Bob Muller and his colleagues have reported that the size of place fields could be modified by increasing the size of the box in which they were recorded but that this scaling was not proportional to the increased area of the box. We wanted to understand this effect in greater detail. In order to do

Place fields stretch as the environment is stretched

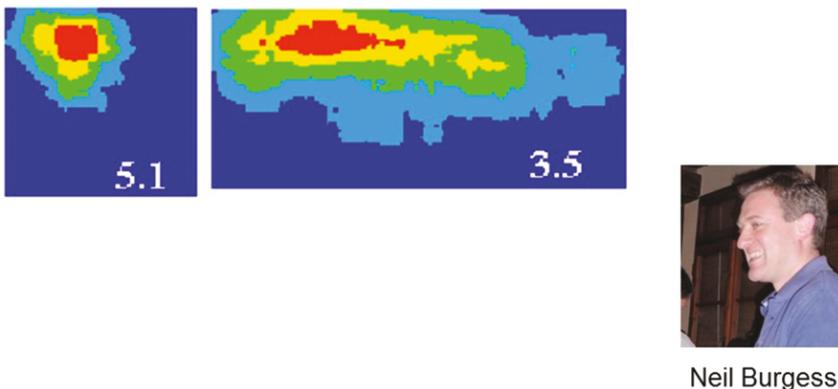


FIGURE 13. Place field stretches along the long dimension when the square is transformed into a rectangle. Number in white is peak firing rate marked by red. After O'Keefe and Burgess (1996).

so, we recorded from place cells in four boxes all placed in the same part of the room but differing in the ratio between the two adjacent sizes of the box. There was a small square, a large square with sides double those of the small square, and two rectangles with adjacent sides equal to the dimensions of the large and small square. The rectangles were essentially small squares stretched out along one dimension. The results show that quite often place fields were stretched along the same dimension as the boxes but were not affected in the orthogonal dimension (Figure 13). Thinking about this effect, Neil and I realised that the simplest model to explain it would be if the inputs to the place cells came from cells which were coding for the distance from a large environmental landmark such as the wall of the box in a particular direction. Cells which stretched along a particular dimension would be attached to the 2 opposing walls. With Tom Hartley, we modelled the firing fields of these cells (Figure 14A) and found that inputs from two or more of these boundary cells when added together followed by thresholding would produce realistic place fields in different environments

Boundary Cells: Theory

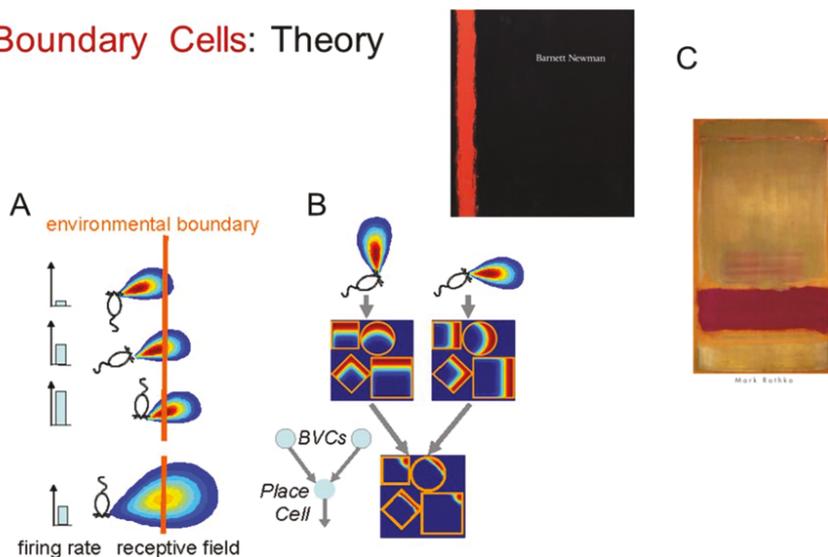


FIGURE 14. Properties of putative boundary cells providing inputs to the place cells to account for the stretching of place fields in the square to rectangle experiment. (A) Firing rate of boundary cells was predicted to increase as the animal got closer to the target distance from the environmental boundaries such as a wall of the box (red line). For some cells the preferred distance from the boundary was shorter (top) than for others (bottom). The fields of boundary cells recorded in different shaped environments were expected to look like the stripes in the paintings of Barnett Newman or Mark Rothko (C). Adding the inputs from 2 or more boundary cells together followed by a thresholding operation would produce place fields. After Hartley et al. (2000).

(B)(Hartley et al., 2000). The putative fields of these boundary cells looked very much like the bold stripes in the paintings of Barnett Newman and Mark Rothko delineating a small part of the two-dimensional space of the painting at a fixed distance from one of the walls (Figure 14C). Boundary cells with the predicted Rothko-like striped fields were subsequently looked for and found (Figure 15A & B) by Colin Lever in our lab (Figure 15 lower right) not in the hippocampus proper but in the neighbouring subiculum (D) (Lever et al., 2009). As predicted by the model, adding the inputs from these recorded two boundary cells followed by thresholding produces a respectable place field (C). Around the same time, similar cells with similar properties were found in the entorhinal cortex in the Moser lab (Solstad et al., 2008). They emphasised the closeness of the fields of many of their cells to the walls of the environment and called them border cells. It is still not clear whether these are the same cell types or the border cells are a subset of the boundary cells.

Head direction cells convey information about the direction in which the animal's head is pointing within the environmental frame (Figure 16A). They were originally identified by Jim Ranck about 1984 in an area called the pre-subiculum (PrS) nestled between the hippocampus and entorhinal cortex and

Boundary Cells in the Subiculum

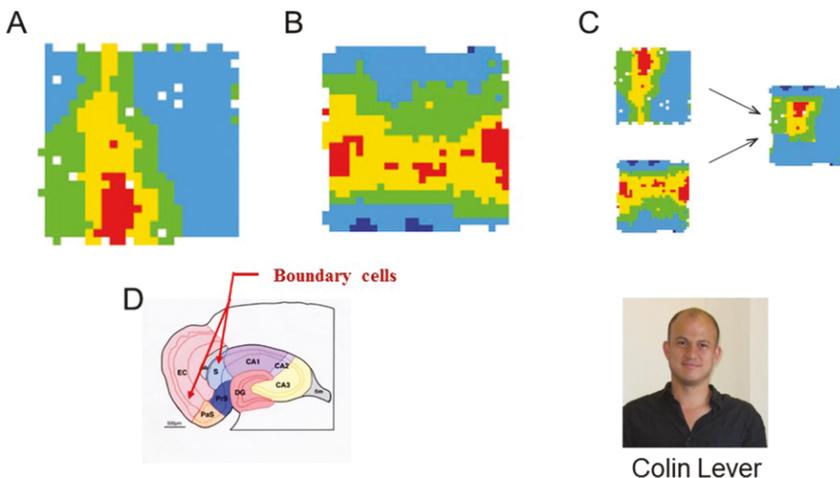


FIGURE 15. Boundary cells were found experimentally in the subiculum (D) by Colin Lever (lower right). Although many had fields close to the walls of the box, a number (A and B) fired away from the boundary and would be expected to produce place-like fields when combined (C). After Lever et al. (2009).

Head Direction Cells

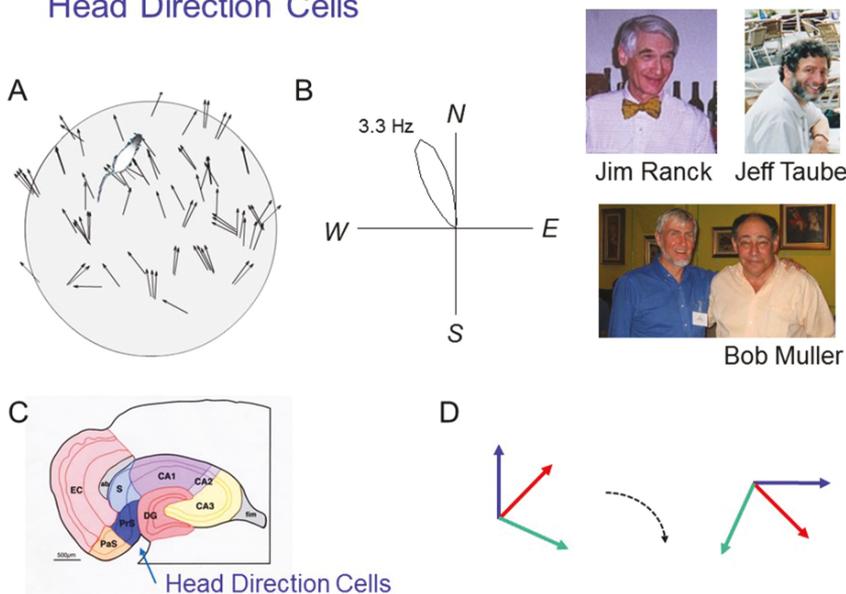


FIGURE 16. Head direction cells fire on the animal's head is pointing in the preferred heading direction (B) irrespective of the animal's location (A). They are found in many parts of the hippocampal circuit including the pre-subiculum (C) where they were discovered by Jim Ranck and studied extensively by Jeff Taube and Bob Muller (upper right). They appear to be wired together (D) such that rotation (dotted arrow) of the preferred direction of one (e.g., blue arrow) is accompanied by an equal rotation of the others (red and green arrows).

were subsequently studied extensively by Jim and his colleagues, Jeff Taube and Bob Muller. Bob Muller, who recently died, had one of the most incisive, original minds of anyone I have had the pleasure to discuss the brain with. The hippocampal spatial field is poorer for his loss. The firing fields of head direction cells are usually represented in a polar coordinate plot as shown in Figure 16B. The cell represented typically fired whenever the animal's head pointed in the North Northwest direction, irrespective of the animal's location in the environment (Figure 16A). An interesting property of these cells is that they seem to be wired together such that their preferred heading directions maintain a constant relationship with each other despite the fact that they can vary in their absolute relationship to the environment. For example, if the animal is briefly disoriented by spinning, the heading direction of a cell which originally pointed in the northward direction might be altered to point in the eastward direction (Figure 16 D, blue arrow); other head direction cells will alter their preferred directions

(red and green arrows) in a concomitant fashion to maintain the internal angles constant. The constellation of cells appears to form a compass-like polar coordinate system upon which the rest of the spatial mapping system is built. As we shall see shortly, head direction cells appear earliest in development, at around the time that the infant rat is leaving the nest and opening its eyes and well before the place and grid components of the map. In addition to inputs about the animal's heading direction, a map needs a distance metric and the question arises as to whether this distance is represented by a one-dimension caliper-like signal in the manner suggested by William Blake's picture of Newton (Figure 17 B) or whether is it something more complicated as shown by the Lichtenstein painting Yellow Cliffs (Figure 17A)? In the latter a grid of dots is laid across the landscape and distance in two directions can be measured from the edge of the figure or the edge of the painting. The answer was provided by research in the Moser lab.

How is distance measured?

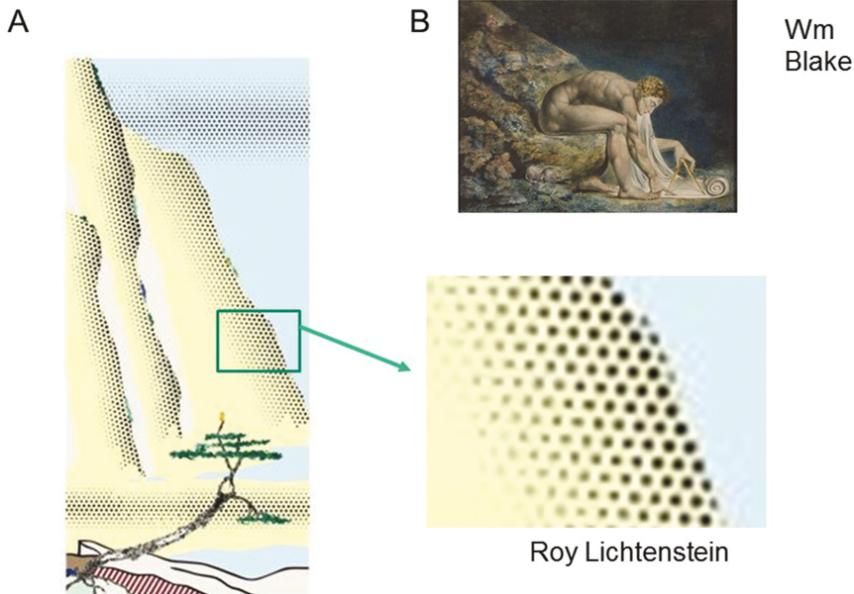
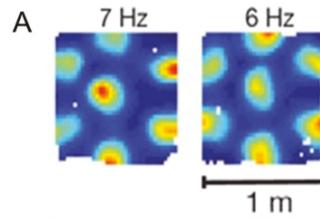
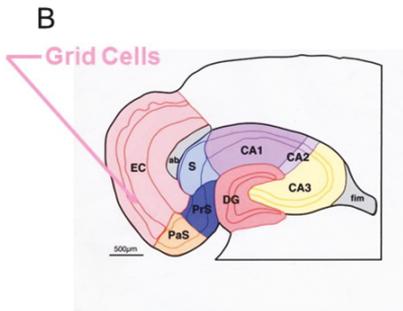


FIGURE 17. Distance within the cognitive map could be measured by a simple one-dimensional calipers as done by William Blake's Newton (B), or by a two-dimensional set of markers laid out on the surface of the environment as in Lichtenstein's painting Yellow Cliffs (A).

Grid Cells: the universal metric in the entorhinal cortex?



May-Britt & Edvard Moser



Torkel Hafting Marianne Fyhn

FIGURE 18. The grid cells (A) of the medial entorhinal cortex (B) discovered by the Torkel Hafting, Marianne Fyhn and the Mosers (lower right) are good candidates for the universal metric of the cognitive map. They fire in a repeating hexagonally symmetrical pattern across the surface of the environment. After Hafting et al. (2005).

Grid Cells are characterised by three variables

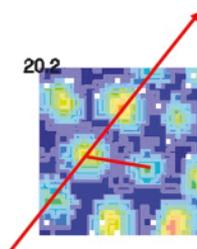


FIGURE 19. The hexagonally symmetrical pattern of grid cell can be described by the distance between two adjacent fields (short horizontal red line), the angle which one of the rows (red arrow) makes with the wall of the box, and the distance of the nearest row from the wall of the box.

Grid Spacings are Quantised

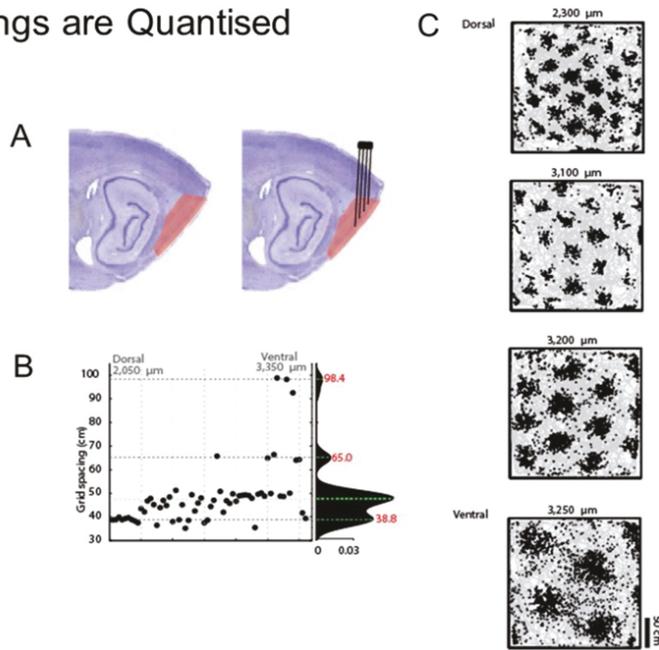
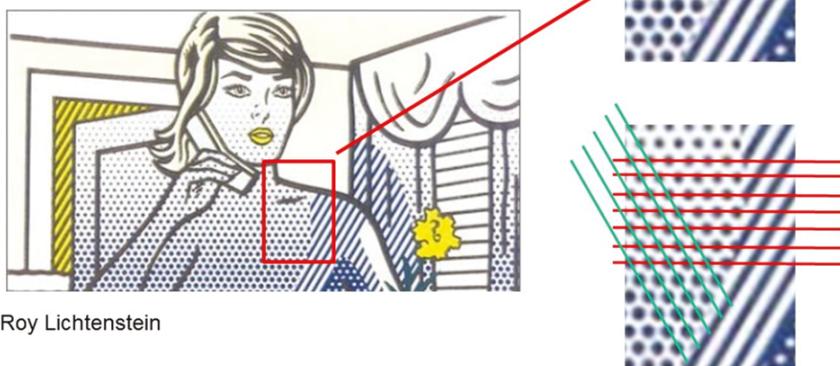


FIGURE 20. The spacings between adjacent peaks of the grid pattern are quantised. Grids more ventral in the entorhinal cortex (A) have a larger metric than ones more dorsal (C) and these changes in metric are quantised, falling into three or four distinct scales (B).

Distance within the mapping system, at least in some environments, could be measured by the entorhinal cortical grid cells (Figure 18). In 2005, the Moser lab reported that there existed in the medial entorhinal cortex (Figure 18B) a set of grid cells which look very much like they represent distances in particular directions: they cover each familiar environment with a set of firing fields which are laid out in a beautiful symmetrical hexagonal pattern and which are reproducible from one trial to next (Figure 18A) (Hafting et al., 2005). It is not hard to see how these patterns could be used to represent distances. As you can see from Figure 19, one can characterise each of these patterns by three variables: the distance between any two adjacent peaks, the smallest offset angle which the orientation of the grid makes relative to the walls and the distance of the grid from the walls of the environment. The cells are grouped into modules with the same orientation and if one looks over a large group of cells in square boxes one finds that they have a preferred orientation relative to the edge of the box with an angle of 8.8° (Krupic et al., 2015). Equally important, the spacing between the peaks of the grid cells varies as a function of location within the entorhinal

cortex. As one goes from dorsal to ventral in the entorhinal cortex (Figure 20 A), scale increases in quantum jumps, each animal having three or four different scales (Figure 20 B,C). This is important because if you take two grid cells with different spacings and you simply add them together and do a little bit of mathematical computation you can produce a place cell. (Figure 24). So we have two ways of producing a place cell, one by the addition of information from two or more boundary cells, and the other by the addition of information from two or more grid cells with different scales. Where do the grid patterns come from? One possibility is that they are created by overlaying different striped patterns as suggested by another painting by Lichtenstein (Figure 21 left). Here he has superimposed his pattern of hexagonal dots over the surface of the painting but in addition juxtaposes to them a series of diagonal stripes (upper right) where you can see that each row of dots is actually the extension of a stripe and you can easily imagine that if you had two more stripes of orientations of 60° degrees relative to each other, the intersection of three stripes patterns would effectively produce the grid pattern (lower right, green and red stripes). Surprisingly, one of our colleagues Julija Krupic (Figure 22 upper right) has found cells which look very much like band cells (A) in the parasubiculum (B), an area which has direct connections to the medial entorhinal cortex.

Relationship between Grids and Stripes

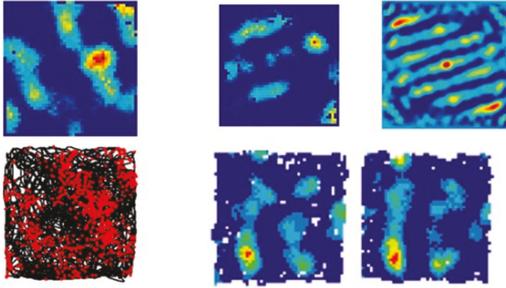


Roy Lichtenstein

FIGURE 21. A simple relationship exists between grids and regular stripes as shown by the section of the Lichtenstein painting expanded in the upper right hand corner where the stripes are aligned with individual rows of the grid pattern. The addition of extra rows of stripes (red and green) at 60° to the original stripes creates the hexagonal pattern.

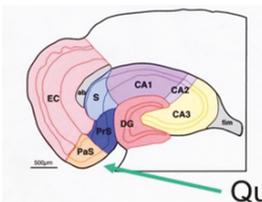
Band-like Cells in the Parasubiculum

A



Julija Krupic

B



Quasi-band cells

FIGURE 22. Band-like cells (A) were found in the parasubiculum (B) by Julija Krupic (upper right). They and other Spatially Periodic Non-grid cells may be members of a large set of hippocampal formation cells which can be decomposed into Fourier components of which the grid cells are the most stable and symmetrical members. After Krupic et al. (2012).

It is very tempting to assume that the regular pattern laid down by the grid cells across different environments might be the basis for the metric of the cognitive map. However we have to be a little cautious about this conclusion, since most of the studies so far have only recorded in circles and squares i.e., highly symmetrical environments. Furthermore the animals were familiar with most of these environments. In recent studies our group has shown that the grid pattern can be modified when the animal is placed in non-symmetrical environment such as a trapezoid (Krupic et al., 2015) and furthermore that there is an expansion of the pattern even in a symmetrical environment such as a square if the environment is novel (Barry et al., 2007). It is clear that we still have a lot to learn about the grid cells and the types of environmental information they are responding to.

The development of the cognitive map provides evidence for some of the ideas of the philosopher Immanuel Kant (Kant, 1963). Kant believed that space along with time was one of the basic organising principles of the human mind and

furthermore that it existed independently and in some sense prior to experience with the objects of the world. He said that “Space is nothing but the form of all appearances of outer sense . . . can be given prior to all actual perceptions, and so exist in the mind *a priori*, and . . . can contain, prior to all experience, principles which determine the relations of these objects” (p. 71). Translated into modern neuroscientific terms, this would predict that some aspects of the cognitive mapping system should be present prior to an animal’s having had any experience with the world. With Francesca Cacucci, Tom Wills and Hui Min Tan (Tan et al., 2015; Wills et al., 2010) (Figure 23 right), we decided to test this by recording from three of the major spatial cell types in very young infant rats in the second, third, and fourth week of life. These animals only begin to leave the nest towards the end of the second week around the same time as their eyes open. What we found is summarised in Figure 23 where one can see that the different cell types emerge at different times. Directional cells are present from P12 just before the time the animal leaves the nest and its eyes are opening. They develop their behavioural correlates very rapidly, quickly approximating adult levels. Some of these directional cells can be recorded prior to significant spatial exploration and eye opening, suggesting their properties might be largely experience-independent. Before eye opening, their fields are unstable and cannot be controlled by the usual visual stimuli. Place cells develop later first appearing at

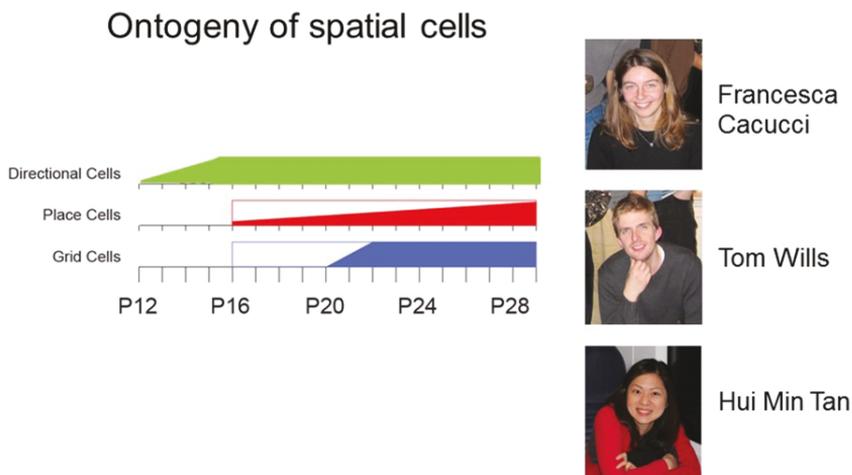


FIGURE 23. Development of spatial cells in the rat. Directional cells appear first as early as postnatal day 12 although they are not stable. They develop rapidly over the next few days to reach adult levels. Place cells first appear around postnatal day 16 and then develop slowly over the next two weeks. Grid cells appear several days later than place cells but develop rapidly. After Wills et al. (2010) and Tan et al. (2015).

Grid cells and Boundary Vector cells may provide 2 independent pathways into Place Representations

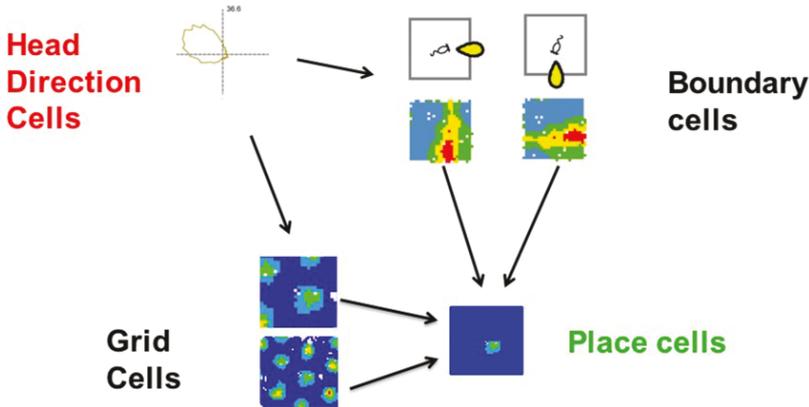


FIGURE 24. Two routes to a place cell. Combinations of boundary cells or grid cells can independently produce a place cell. Both boundary cells and grid cells appear to receive inputs from head direction cells.

P16, and increasing in number and spatiality during the next two weeks. Last to develop are the grid cells, which only appear in significant numbers at P22 but then quickly reach adult levels. Similar findings were published at the same time by the Moser lab (Langston et al., 2010). It appears that one aspect of the cognitive map, the directional system, develops prior to any substantial relevant experience with the environment and it probably provides inputs to the place and grid cells forming the basic substrate for the rest of the cognitive map. The fact that the place cells develop before the grid cells is further support for the idea that there are two or more pathways to create a place cell only of one which is dependent on the grid cell input. Figure 24 shows how place cells can be created either by the intersection of inputs from two boundary cells with stripes oriented at an angle to each other or by the intersection of inputs from two grid cells with different spacings.

Virtual reality technology can be used to test this idea that there are two different ways in which place cells can be generated (Chen et al., 2013). Our setup created by Guifen Chen and John King (Figure 25 right) derives from work originating in Tübingen and perfected in Princeton by David Tank's lab. The animals run on a Styrofoam ball which floats on a cushion of air and the motion of the ball and thus the animal is tracked by a computer mouse (Figure 25 left). Figure 25 (mid right) shows a screenshot of the linear track which is projected onto the

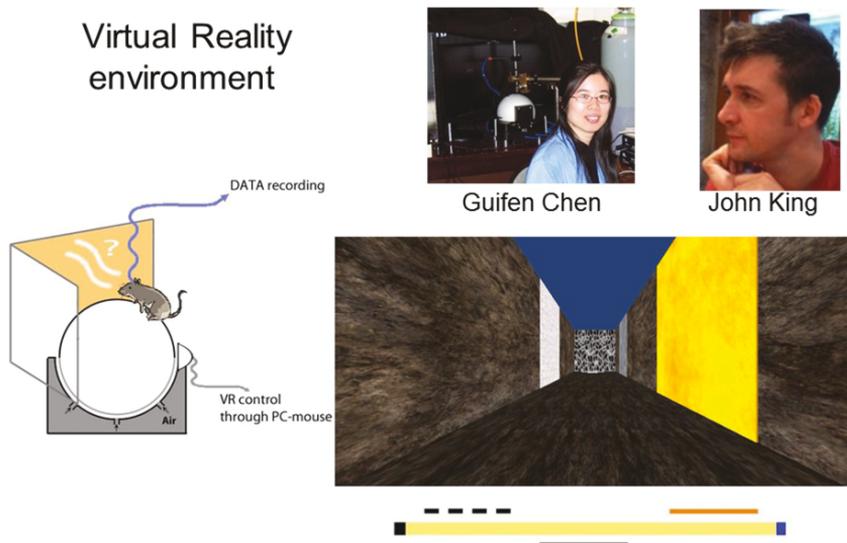


FIGURE 25. A Virtual Reality environment for mice was created by Guifen Chen and John King (upper right) and consists of a Styrofoam ball on top of which the animal runs and where his movements are translated into movements in a virtual environment projected onto VDU screens. A screen shot from the virtual linear track (middle right) shows the visual cues on the ends and sides of the track. The track and visual cues are schematised by the long yellow bar and coloured stripes (below).

two video displays in front of the animal. The linear track and the visual cues are schematically represented as shown in the lower right hand of the figure. The firing fields of 80% of the cells are dependent on the visual cues on the walls of the virtual environment as shown by the loss of the field when these cues are removed (Figure 26A middle). Interestingly it is the side cues which are important and not the cues on the ends of the track (bottom). The importance of the visual cues is confirmed by the maintenance of the place field in 25% of the cells on trials in which the animal is passively moved through the environment at about the same speed as it normally runs (Figure 26B middle). One way in which the path integration system can control the place fields is shown in Figure 27 where the animal was started at the beginning of the track with the visual cues present to allow it to locate itself but then these were turned off as the rat started to move down the track and were turned on again as it reach the goal (Figure 27 bottom). In almost half of the cells the fields were maintained during this procedure showing that once the animal had oriented itself in an environment, information from its own movements was sufficient to update the representation and cause the cell to fire in the correct location. The visual orientation at

Control by visual cues on the side wall

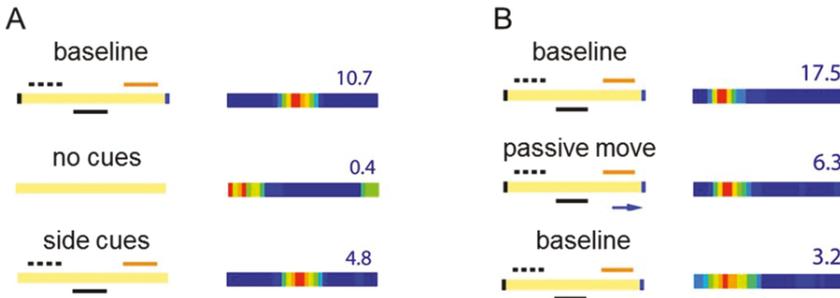


FIGURE 26. Place cells recorded in the virtual environment depend on the visual cues particularly those on the side walls. A, a place cell which stops firing when all the visual cues are removed (middle row) but continued to fire if only the cues on the end walls are removed (bottom row). B, a place cell which continues to fire if the animal is passively moved along the track. After Chen et al. (2013).

Path integration (cues off trial)

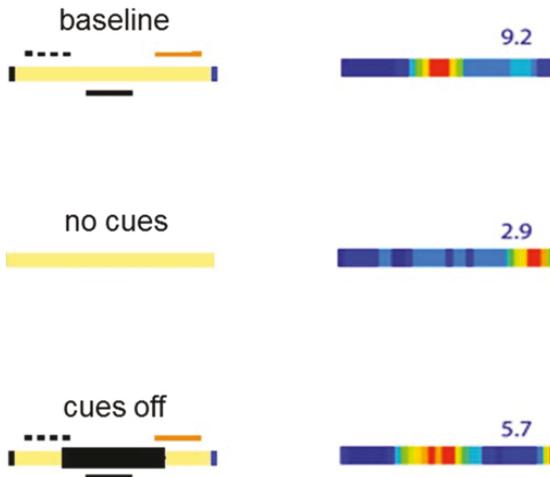


FIGURE 27. Place cells can combine visual and path-integration cues. The cell does not fire if the visual cues are removed throughout the trial (middle row) but does fire appropriately if they are available at the beginning of the run and then turned off when the animal runs through the field (bottom row). After Chen et al. (2013).

the beginning of the trial was necessary, as shown by the fact that the cells did not fire in the correct location if the visual cues were removed throughout the trial (Figure 27 middle).

The human hippocampus is widely believed to store episodic memories and the question arises as to whether it also involved in spatial memory and navigation similar to the rat. It is more difficult to study navigation in humans than in small animals and it was only when virtual reality technology became cheap and widely available in the 1990s that we were in a position to do so. That was when first person video games became available and more importantly came with editing facilities which allowed modifications to be made to the environment. Neil Burgess modified a game called Duke Nukem by removing the guns and monsters leaving us with a 70×70 metre environment containing many different rooms (Figure 28) and importantly many different routes by which the participant could move from one location to another (Figure 29 B). Eleanor Maguire (Figure 28 right) and Neil carried out a study in which they allowed normal volunteers to become familiar with this environment and scanned their brains as they found their way around this environment using either a cognitive

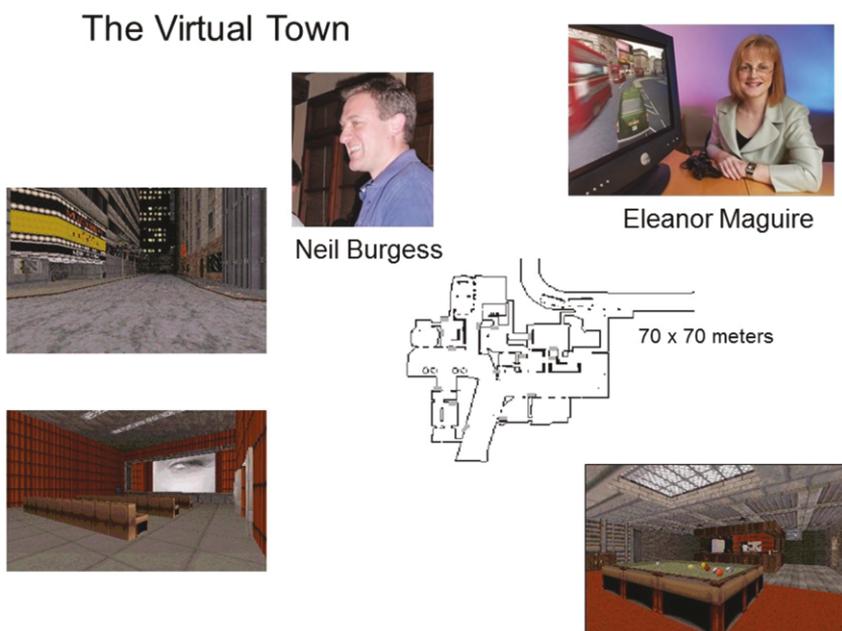


FIGURE 28. Screenshots from the virtual town created by Neil Burgess and used by himself and Eleanor Maguire to test hippocampal function in humans. The centre panel shows the layout of the environment.

map strategy or the route based strategy of following a series of arrows on the floor (Maguire et al., 1998). When we looked at the differences between the activations in these conditions we found the right hippocampus was more active in the cognitive map strategy (Figure 29A). Importantly the amount of activity in the hippocampus increased as a function of the accuracy with which the participant moved from one location to another (Figure 29 C) suggesting that the more active the hippocampus, the better the navigation. Maguire and colleagues (Maguire et al., 2000) then went on to show that London taxicab drivers had a larger posterior hippocampus than controls (Figure 30 A) and that this increase in size was directly related to the amount of time as a cab driver showing that the important factor was the experience of London's complicated street patterns and that London cab drivers are probably not born with bigger hippocampi but develop them.

Hippocampal Activation in Map-Based navigation

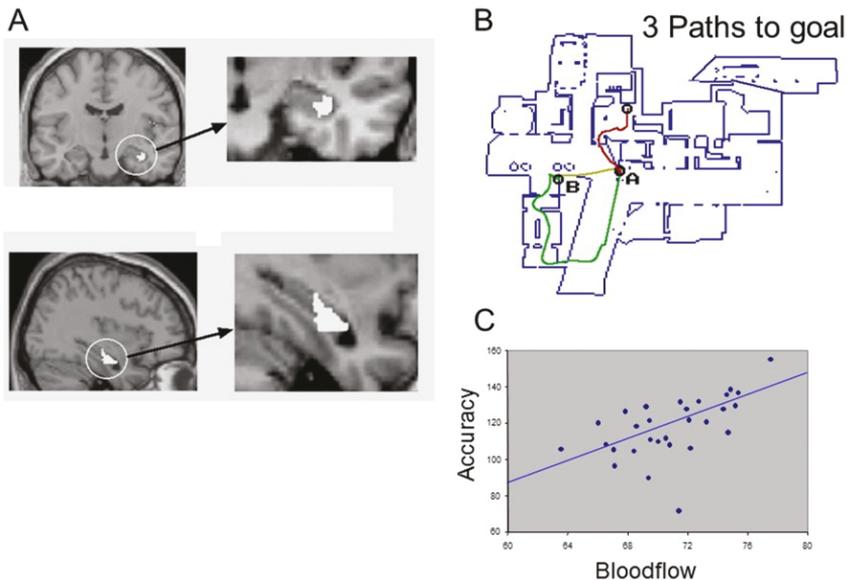


FIGURE 29. Right hippocampus is activated in map-based navigation. PET scans (A) show activation in the right hippocampus when participants used cognitive mapping strategies to find their way around the environment rather than route based ones. More accurate navigation (e.g., yellow line from A to B in panel B) is correlated with more bloodflow in the hippocampus than more circuitous or ineffective routes (green and red lines) (C). After Maguire et al. (1998).

Posterior Hippocampus is LARGER in taxicab drivers and increases with experience

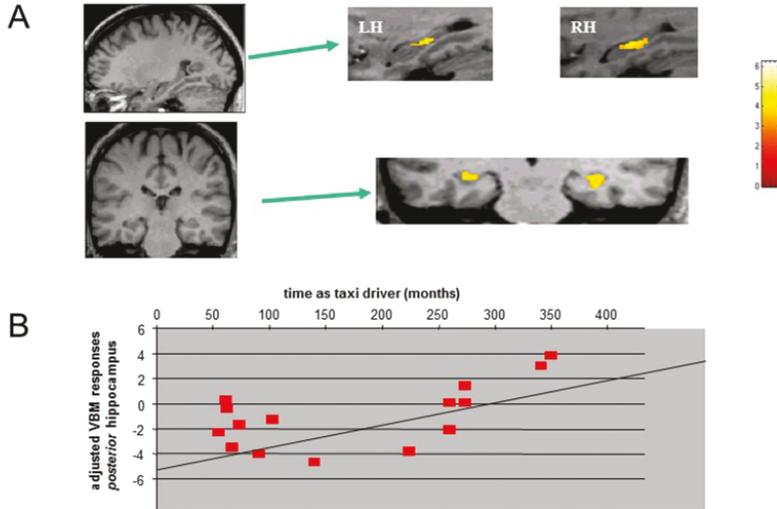


FIGURE 30. The posterior hippocampus is larger in taxicab drivers (A) than in controls and this effect increases with the length of time as a cab driver (B). After Maguire et al., 2000).

In Summary, the Hippocampal Formation provides a cognitive map of a familiar environment which can be used to identify the animal's current location and to navigate from one place to another. The Mapping system consists of a number of different spatial cells including ones which identify the animals place, direction, distance from landmarks such as the walls of an environment in a particular direction (boundary cells) and in some environments a metric for measuring distances between points in the map (grid cells). The grid cells cannot provide a perfect metric in all environments since the grid pattern is distorted in asymmetrical environment such as trapezoids and the scale of the grids is increased in unfamiliar environment. There is evidence for 2 independent strategies for locating places, one based on environmental landmarks probably provided by the boundary cells and the other on a path integration system which uses information about distances travelled in particular directions provided by the grid cells.

A similar spatial system exists in humans which additionally provides the basis for human episodic memory. A prototypical episodic memory system

could be achieved by the addition of a sense of linear time to the basic spatial system seen in the rat (Burgess et al., 2002).

REFERENCES

1. Barry C, Hayman R, Burgess N, Jeffery KJ. (2007). Experience-dependent rescaling of entorhinal grids. *Nat Neurosci* **10**:682–684.
2. Burgess N, Maguire EA, O'Keefe J. (2002). The human hippocampus and spatial and episodic memory. *Neuron* **35**:625–641.
3. Chen G, King JA, Burgess N, O'Keefe J. (2013). How vision and movement combine in the hippocampal place code. *Proc Natl Acad Sci U S A* **110**:378–383.
4. Corkin S. (2013). *Permanent present tense*. Basic books.
5. Hafting T, Fyhn M, Molden S, Moser MB, Moser EI. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**:801–806.
6. Hartley T, Burgess N, Lever C, Cacucci F, O'Keefe J. (2000). Modeling place fields in terms of the cortical inputs to the hippocampus. *Hippocampus* **10**:369–379.
7. Huxter J, Burgess N, O'Keefe J. (2003). Independent rate and temporal coding in hippocampal pyramidal cells. *Nature* **425**:828–832.
8. Kant, Immanuel. (1963). *Critique of pure reason*. (Translated by N Kemp Smith—original published 1787). London: Macmillan.
9. Krupic J, Bauza M, Burton S, Barry C, O'Keefe J. (2015). Grid cell symmetry is shaped by environmental geometry. *Nature* **518**:232–235.
10. Krupic J, Burgess N, O'Keefe J. (2012). Neural representations of location composed of spatially periodic bands. *Science* **337**:853–857.
11. Langston RF, Ainge JA, Couey JJ, Canto CB, Bjerknes TL, Witter MP, Moser EI, Moser MB. (2010). Development of the spatial representation system in the rat. *Science* **328**:1576–1580.
12. Lever C, Burton S, Jeewajee A, O'Keefe J, Burgess N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *J Neurosci* **29**:9771–9777.
13. Maguire EA, Burgess N, Donnett JG, Frackowiak RS, Frith CD, O'Keefe J. (1998). Knowing where and getting there: a human navigation network. *Science* **280**:921–924.
14. Maguire EA, Gadian DG, Johnsrude IS, Good CD, Ashburner J, Frackowiak RS, Frith CD. (2000). Navigation-related structural change in the hippocampi of taxi drivers. *Proc Natl Acad Sci U S A*.
15. McNaughton BL, Barnes CA, O'Keefe J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp Brain Res* **52**:41–49.
16. Morris RGM. (1981). Spatial localisation does not depend on the presence of local cues. *Learning and Motivation* **12**:239–260.
17. O'Keefe J. (1996). The spatial prepositions in English, vector grammar and the cognitive map theory. In: *Language and Space* (Bloom P, Peterson M, Nadel L, Garrett M, eds), pp. 277–316. Cambridge, Mass.: MIT Press.

18. O'Keefe J, Nadel L. (1979). Precis of O'Keefe and Nadel's The hippocampus as a cognitive map. *The Behavioral and Brain Sciences* **2**:487–533.
19. O'Keefe J, Nadel L. (1978). *The hippocampus as a cognitive map*. Oxford University Press.
20. O'Keefe J. (1976). Place units in the hippocampus of the freely moving rat. *Exp Neurol* **51**:78–109.
21. O'Keefe J, Burgess N, Donnett JG, Jeffery KJ, Maguire EA. (1998). Place cells, navigational accuracy, and the human hippocampus. *Philos Trans R Soc Lond B Biol Sci* **353**:1333–1340.
22. O'Keefe J, Dostrovsky J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* **34**:171–175.
23. O'Keefe J, Recce ML. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* **3**:317–330.
24. O'Keefe J, Speakman A. (1987). Single unit activity in the rat hippocampus during a spatial memory task. *Exp Brain Res* **68**:1–27.
25. Scoville WB, Milner B. (1957). Loss of recent memory after bilateral hippocampal lesions. *J Neurol Neurosurg Psychiatry* **20**:11–21.
26. Solstad T, Boccara CN, Kropff E, Moser MB, Moser EI. (2008). Representation of geometric borders in the entorhinal cortex. *Science* **322**:1865–1868.
27. Tan HM, Bassett JP, O'Keefe J, Cacucci F, Wills TJ. (2015). The development of the Head Direction System before Eye Opening in the Rat. *Curr Biol* **25**:479–483.
28. Tolman EC. (1948). Cognitive maps in rats and men. *Psychol Rev* **55**:189–208.
29. Vanderwolf CH. (1969). Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalogr Clin Neurophysiol* **26**:407–418.
30. Wills TJ, Cacucci F, Burgess N, O'Keefe J. (2010). Development of the hippocampal cognitive map in preweanling rats. *Science* **328**:1573–1576.
31. Yartsev MM, Ulanovsky N. (2013). Representation of three-dimensional space in the hippocampus of flying bats. *Science* **340**:367–372.

Portrait photo of John O'Keefe by photographer Alexander Mahmoud.