

A metric for the cognitive map: found at last?

Kathryn J. Jeffery¹ and Neil Burgess²

¹Institute of Behavioural Neuroscience, Department of Psychology, University College London, 26 Bedford Way, London WC1H 0AP, UK

²Institute of Cognitive Neuroscience and Department of Anatomy, University College London, 17 Queen Square, London WC1N 3AR, UK

A network of brain areas collectively represent location, but the underlying nature of this ‘cognitive map’ has remained elusive. A recent study reports that the activity patterns of some entorhinal cortical neurons form a remarkably regular array of evenly spaced peaks across the surface of the environment. These ‘grid cells’ might be the basis of a metric used for calculating position, and their discovery could greatly advance our understanding of how navigational computations are performed.

An enduring debate in spatial cognition concerns whether the brain generates a truly map-like representation of the environment, with an intrinsic metric structure, or a looser more topological representation of environmental features and their relationships. This debate goes back at least to the middle of the 20th century (e.g. [1]) and was rekindled by the discovery in the 1970s of neurons in the hippocampus of rats that exhibited spatially localized activity [2] (see Box 1). These ‘place cells’ (Figure 1a) were the first pointer to a role for the hippocampus and surrounding areas in spatial representation, and again the question arises: does each place cell respond to conjunctions of sensory features unique to a given place, or does some underlying metric enable feature-independent estimation of location? Now, Hafting *et al.* [3] have discovered ‘grid cells’, one synapse upstream of the place cells, in medial entorhinal cortex (mEC). Each grid cell shows multiple peaks of firing distributed regularly across the surface of the rat’s environment (Figure 1b). The constant spacing of the peaks for a given cell could not plausibly be generated by individual associations to environmental features, and strongly suggests the presence of an intrinsic metric. This metric grid could enable the spatial computations underlying our sense of place. To complement the ‘place’ sense, there are also head-direction cells found in numerous structures within the limbic circuit, which provide a sense of direction (Figure 1c).

Properties of the entorhinal grids

Hafting *et al.* recorded from entorhinal neurons as rats foraged in various enclosures, and these neurons proved to have highly organized firing patterns, each neuron showing multiple activity peaks arranged in a hexagonally close-packed triangular grid covering the environment

(Figure 1b). The inter-peak distance was constant for a given cell and for a given anatomical location, but steadily increased dorsoventrally within mEC, from around 30 cm to around 60 cm (Figure 2b). The orientation of the grids (Figure 2c) in a particular environment was also constant for a given cell, and was similar among closely neighbouring cells, but varied across the broader population. Although aligned with each other, the grids of closely neighbouring cells had different relative offsets, so that small sets of neighbouring grid cells could effectively carpet the whole environment (Figure 2d).

These remarkable findings strongly imply a neural metric, analogous to multiple sheets of graph paper of varying orientations and scales, which might enable the spatial computations required to determine locations and translations.

Extrinsic and intrinsic influences on grid cells

The influence of external sensory inputs on grid-cell firing is indicated by the reproducible location and orientation of grids across recording sessions. In addition, rotation of a single directional landmark caused a corresponding rotation of the grid. Notwithstanding these external influences, the most important aspect of this discovery is its implication of an internal metric, allowing a given cell to space its firing peaks evenly regardless of the sensory environment – as confirmed by the resistance of grid scale to changes in the size of the environment (Figure 2a).

What could be the source of a grid cell’s intrinsic metric? To maintain a fixed inter-peak distance, the cells probably use some kind of movement-based information. Movement-based cues are known to contribute both to place-cell localization and to animal navigation more generally. Speed and direction signals can be combined to estimate translation, a process known as ‘path integration’ [4], and understanding how grid cells combine such signals could shed light on the neural mechanisms underlying path integration.

Functional significance

The grid cells raise several intriguing questions, most obviously, what are they for? McNaughton *et al.* [5] suggested that the hippocampal navigational system evolved to represent spatial relationships without explicit reference to objects, the system being updated on the basis of the animal’s movements [2,5,6]. The entorhinal cortex, it now seems, might contain just such a featureless representation: a purely metric grid, not associated with

Corresponding author: Jeffery, K.J. (k.jeffery@ucl.ac.uk).

Available online 23 November 2005

Box 1. Building a cognitive map

The discovery of 'place cells' (see Figure 1a in main text) in the rat hippocampus gave rise to the notion that the limbic circuitry and surrounding regions might cooperate in forming a 'cognitive map' of the environment [2]. In an open environment, place cells fire whenever the animal enters a specific region, independently of the animal's orientation or the availability of any particular subset of external cues. Cells with similar properties have been reported in monkeys [12,13] and humans [14]. The grid cells, which are located one synapse upstream of the place cells, appear to provide a metric on which to base such a map. What are the other components?

Spatial navigation requires a sense of direction as well as of place, and this could be provided by the head-direction cells found in numerous structures within the limbic circuit [15] (Figure 1c in main text). Each head-direction cell becomes active when the animal faces in a particular direction, irrespective of its location. The orientation of

head-direction cells is maintained by both landmark information and the animal's sense of self-motion. Head-direction cells in the pre-subiculum, which projects to medial entorhinal cortex, might provide the information needed to orient the grid arrays.

A map also needs something to identify it (akin to a page number, or a title), in order to define which environment a given co-ordinate refers to. In the hippocampus, this could be supplied by the subset of active place cells that is unique to a given environment. When the animal leaves one environment and enters a different one, the set of place cells 'remaps': it takes on a different firing pattern, triggered by the different environmental features [10]. Although data are still sparse at present, it is possible that grid cells might also remap, in the sense of expressing altered offsets in different environments. If confirmed by future experiments, this could provide a possible basis for place cell remapping (see [8]).

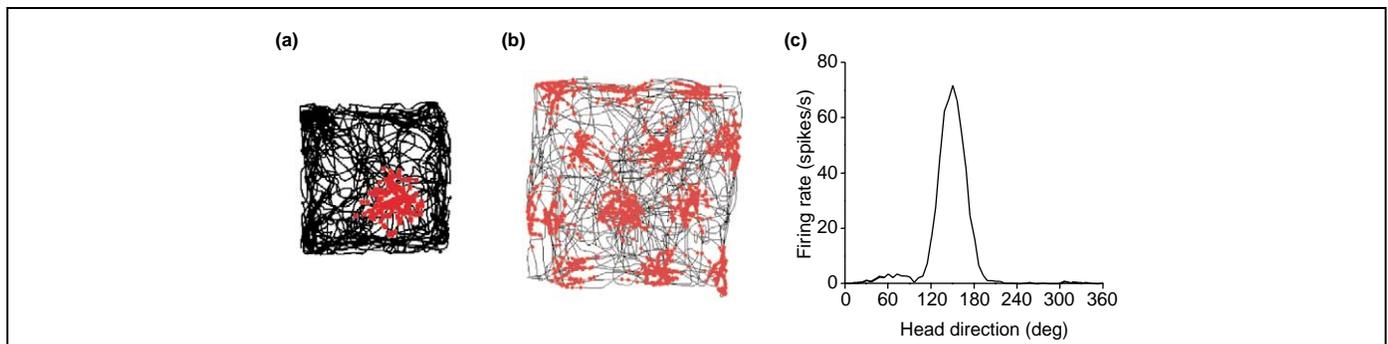


Figure 1. (a) A typical 'place field' generated by a hippocampal neuron. The data were obtained by recording hippocampal neurons as a rat foraged in a 70×70 cm square enclosure. The black tracing shows the path of the rat during the 8-min trial, and the red dots show the location of the rat when each action potential of the place cell fired. As is typical of place cells, the cell only fired when the rat entered a particular region of the environment, and was silent outside that region. (b) Examples of the activity of an entorhinal 'grid cell' recorded from the dorsal medial entorhinal cortex as the rat foraged in a 1 m square enclosure; lines and red dots indicate the same as in (a). The cell fired in spatially restricted regions, as with place cells, but firing regions are repeated at regular intervals across the environment to produce a multi-peaked triangular array. (Adapted from [3].) (c) The activity of a head-direction cell. Note how the cell's activity increased sharply to a peak when the rat's head was within a restricted range of directions. Despite this directional restriction, such cells show activity at all locations across the environment. (Figure kindly supplied by Dr Jeffrey Taube.)

any environment in particular but useful for navigating them all. Hafting *et al.* note that the regular pattern of grid-cell firing might enable a path-integration process because small sets of grid cells can cover an arbitrarily

large area (see also [7]). In addition, the constant displacement of sets of neighbouring grids indicates that the interconnections within such a set of grid cells could be endlessly tuned to implement the effects of self-motion

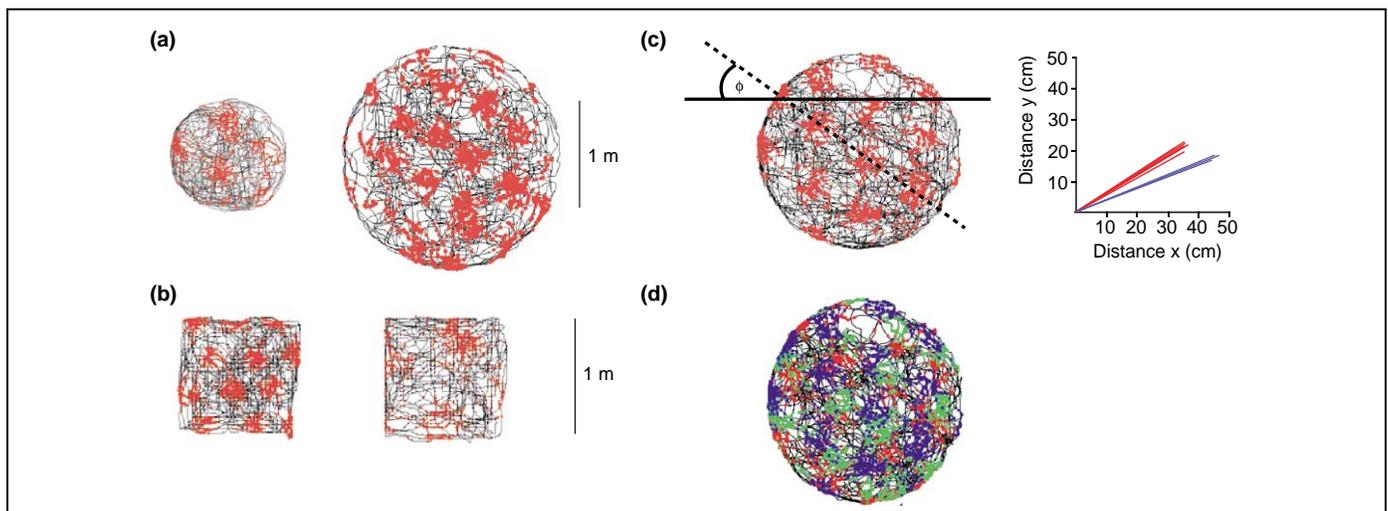


Figure 2. (a) Resistance of the grid metric to changes in environment scale. Note that the distance between firing peaks was similar whether the rat was in a small environment (left) or a larger one (right). (b) Conversely, the interpeak distance varied according to the recording location within medial entorhinal cortex: closer spacing more dorsally (left) vs. wider spacing more ventrally (right). (c) Left: the orientation of the grids was assessed by drawing an imaginary line through a row of peaks and comparing this with a camera-defined 'horizontal'. Cells that were close together tended to have similar orientations. Right: vector illustration of orientation (shown by the orientation of the coloured lines) showing that dorsally located cells (red) tended to have smaller grid spacings than those recorded more ventrally (blue), and that cells within one location all tended to have similar grid orientation. (d) When the firing fields of three nearby grid cells were superimposed, they covered most of the environment with aligned but slightly offset grids. (Adapted from [3].)

accurately, irrespective of the animal's location. By contrast, place cells might mediate the association of grids to the environment, allowing the selection of grid cells that overlap at a given location to be associated with the specific sensory input available at that location via connections with a specific place cell [8].

The finding of grid cells opens up many avenues of enquiry, for example, how do grid cells and place cells together contribute to phenomena such as remapping (see Box 1), or the encoding of environmental shape [9,10]? What implications do grid cells have for the temporal (e.g. [8]) and non-spatial (e.g. [11]) processing of information in this brain region? Most importantly, the discovery of grid cells serves to deepen further the question of how these brain structures contribute to episodic memory, which remains the most obvious and yet mysterious function of this network in humans.

References

- 1 Tolman, E.C. (1948) Cognitive maps in rats and men. *Psychol. Rev.* 40, 40–60
- 2 O'Keefe, J. and Nadel, L. (1978) *The Hippocampus as a Cognitive Map*, Clarendon Press
- 3 Hafting, T. *et al.* (2005) Microstructure of a spatial map in the entorhinal cortex. *Nature* 436, 801–806
- 4 Etienne, A.S. and Jeffery, K.J. (2004) Path integration in mammals. *Hippocampus* 14, 180–192
- 5 McNaughton, B.L. *et al.* (1996) Deciphering the hippocampal polyglot: the hippocampus as a path integration system. *J. Exp. Biol.* 199, 173–185
- 6 Zhang, K. (1996) Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J. Neurosci.* 16, 2112–2126
- 7 Hargreaves, E.L. *et al.* (2005) Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. *Science* 308, 1792–1794
- 8 O'Keefe, J. and Burgess, N. (2005) Dual phase and rate coding in hippocampal place cells: theoretical significance and relationship to entorhinal grid cells. *Hippocampus* 15, 853–866
- 9 O'Keefe, J. and Burgess, N. (1996) Geometric determinants of the place fields of hippocampal neurons. *Nature* 381, 425–428
- 10 Muller, R.U. and Kubie, J.L. (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci.* 7, 1951–1968
- 11 Wood, E.R. *et al.* (1999) The global record of memory in hippocampal neuronal activity. *Nature* 397, 613–616
- 12 Hori, E. *et al.* (2003) Representation of place by monkey hippocampal neurons in real and virtual translocation. *Hippocampus* 13, 190–196
- 13 Rolls, E.T. (1999) Spatial view cells and the representation of place in the primate hippocampus. *Hippocampus* 9, 467–480
- 14 Ekstrom, A.D. *et al.* (2003) Cellular networks underlying human spatial navigation. *Nature* 425, 184–188
- 15 Taube, J.S. (1998) Head direction cells and the neurophysiological basis for a sense of direction. *Prog. Neurobiol.* 55, 225–256

1364-6613/\$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.tics.2005.11.002

Letters

Does the normal brain have a theory of mind?

Valerie E. Stone¹ and Philip Gerrans²

¹School of Psychology, University of Queensland, McElwain Bldg, St Lucia, QLD 4072, Australia

²Department of Philosophy, University of Adelaide, Australia

Apperly, Samson and Humphreys [1] have elegantly detailed the empirical standards necessary to claim that theory of mind (ToM) is domain-specific. They argue convincingly that current evidence from research with neurological patients under-determines domain-specific claims, but do not describe the alternative to the domain-specific view. We sketch an explicit alternative, describing computational architecture that could support ToM inferences without requiring a specific ToM module. We argue that this view integrates evidence from both autism and neuropsychology more convincingly than the modular view.

ToM abilities depend on the interaction (both developmental and on-line) of domain-general abilities with lower-level cognitive mechanisms for representing social information: face processing, gaze monitoring, tracking of intentions and goals, and joint attention [2–5]. These lower-level mechanisms are *domain-specific* – restricted to social stimuli and dependent on specific neural circuitry [5]. Their normal functioning is an essential precursor to normal ToM performance [2,4,5]. However, they are not sufficient by themselves for sophisticated ToM (belief-

state) inferences. The outputs of these lower-level mechanisms are used for inferences by higher level *domain-general* mechanisms: executive function, metarepresentation and recursion [4–7]. Executive function allows us to keep the elements of a social interaction in mind, and inhibit our own knowledge of the state of reality when asked about someone else's mental state [5]. Metarepresentation operating on information about eye gaze and attention (who saw or was attending to what) allows us to represent others' knowledge states (who knew what) [5]. Recursion operating on metarepresentations of mental states allows us to reason about not just others' thoughts, but others' thoughts about thoughts [7].

On our view, ToM is no more than what happens when these domain-general mechanisms interact with lower-level, domain-specific mechanisms to process social information [4,5]. Deficits on ToM tasks can result from deficits in low-level social input systems (e.g. joint attention) *or* in higher-level domain-general capacities. On this view, it should be impossible to find a pure ToM deficit occurring independently of other deficits.

Indeed, there is currently no evidence for a pure ToM deficit. Children with autism have deficits not only on ToM

Corresponding author: Stone, V.E. (stone@psy.uq.edu.au).

Available online 28 November 2005