

Published in final edited form as:

Q J Exp Psychol (Hove). 2010 September ; 63(9): 1683–1699. doi:10.1080/17470210903540771.

Theoretical accounts of spatial learning – a neurobiological view (commentary on Pearce 2009)

Kathryn J Jeffery

Institute of Behavioural Neuroscience Department of Cognitive, Perceptual and Brain Sciences
Division of Psychology and Language Sciences University College London 26 Bedford Way
London WC1H 0AP

Abstract

Theories of learning have historically taken, as their starting point, the assumption that learning processes have universal applicability. This position has been argued on grounds of parsimony, but has received two significant challenges: first, from the observation that some kinds of learning, such as spatial learning, seem to obey different rules from others, and second, that some kinds of learning take place in processing modules that are separate from each other. These challenges arose in the behavioural literature but have since received considerable support from neurobiological studies, particularly single neuron studies of spatial learning, confirming that there are indeed separable (albeit highly intercommunicating) processing modules in the brain which may not always interact (within or between themselves) according to classic associative principles. On the basis of these neurobiological data, reviewed here, it is argued that rather than assuming universality of associative rules, it is more parsimonious to assume sets of locally operating rules, each specialised for a particular domain. By this view, although almost all learning is associative in one way or another, the behavioural-level characterization of the rules governing learning may vary depending on which neural modules are involved in a given behavior. Neurobiological studies, in tandem with behavioural studies, can help reveal the nature of these modules and the local rules by which they interact.

Keywords

associative learning; spatial learning; geometric module; place cells; grid cells

Introduction

Associative learning theory, which began with the seminal studies of Pavlov, is perhaps the most complete account of animal learning that exists today. It rests on the proposition that learning involves the formation of associations between representations, and the classical view has been that it is guided by specific rules that operate in a domain-independent manner (that is, more or less irrespective of the nature of the information being represented or of the sensory modalities involved). These associative rules mostly pertain to how different environmental cues compete with each other for influence over the animal's behaviour, and the theory has been developed over several decades into a formalism that is successful in explaining many otherwise counterintuitive behavioural results.

The initial position of learning theorists was that learning rules are likely to be universal, operating independently of the knowledge domain (be it spatial map, object representation, etc). More recently, however, theorists have started to recognize that behavioural data are better explained by models in which multiple rules interact in different ways to specify learning in different domains (e.g., Le Pelley, 2004). The purpose of this article is to show how neurobiological evidence supports this quasi-modular view of learning, and suggests that different modules operate on different kinds of information, possibly by different rules. Neurobiological studies can help determine what the modules are and the rules by which they interact locally, and these data can help the formulation of behavioral-level models of learning.

The problem of spatial learning

Research into learning began with conditioning studies in which learning occurs in the time domain (e.g., learning that a stimulus predicts a subsequent response). When studies of spatial learning, initiated by Tolman in the 1930s and 40s, began to gain momentum in the 1960s and thereafter, learning theorists assumed that the principles elucidated from conditioning studies would probably turn out to be true for spatial processing as well. This view fairly quickly became challenged from two directions. First, cognitive map theory, first articulated by Tolman (Tolman, 1948) and then developed and formalised by O’Keefe and Nadel (O’Keefe & Nadel, 1978), suggested that the principles governing spatial behaviour may not be the same as those governing Pavlovian and operant conditioning behaviour in Skinner boxes. Specifically, O’Keefe and Nadel argued that spatial processing seems to operate automatically, even in the absence of reinforcement, and also seems to be domain specific such that different kinds of information are processed differently: including, perhaps, according to different “rules”. Furthermore, relational qualities such as distance and adjacency appear to be explicitly extracted and represented in the brain, in the kind of “black box” processing that was disallowed by the original behaviourist formulation of learning. Second, Cheng and his colleagues suggested that spatial processing involves, first and foremost, the operation of a “geometric module” in which a relatively autonomous brain area operates specifically on geometric information – again, suggesting domain specificity (Cheng, 1986; Cheng & Newcombe, 2005); but see (Cheng, 2008b) for an updated view).

Both of these challenges to associative learning theory have inspired a substantial body of research attempting to decide between alternative accounts of spatial behaviour. One of the arguments about spatial learning and the cognitive map, detailed below, has concerned whether phenomena in the spatial domain do indeed always conform to the principles of associative learning. The debate about the geometric module has centred on whether geometric information has some kind of privileged role within the spatial representation, or whether geometric cues are just another form of cue. The question has also arisen as to whether geometry is, as it were, “impervious” to other sorts of information (such as landmarks, or the colour and texture of walls). These two lines of argument, about associative rules and modularity, come together in a number of recent experiments that have explored whether associative learning phenomena can account for cue integration, with and without geometry, in the spatial domain. The unresolved question is whether these findings in the spatial domain render the universal-rules position no longer tenable.

Pearce (2009) has offered a lucid review of an elegant set of experiments, conducted over several years, to explore these issues by examining the ways in which animals use environmental cues to guide spatial behaviour. Some of these studies show that cue processing obeys classic associative learning rules, and the animals’ behaviour can thus be explained by the same rules governing operant conditioning, without recourse to less parsimonious notions such as maps or modules. However, he notes that these rules do not seem to be universally applicable, and that the deviation from these rules occurs with certain

kinds of cue interaction. Pearce suggests that further experiments are needed to evaluate if boundaries and/or the shape of the environment are different in some way, and to determine the nature of the local cues that guide geometry-based navigation. Implicit in this conclusion is the possibility that not all kinds of cue interaction operate according to the same rules.

This possibility seems eminently plausible when one considers the underlying machinery of spatial learning. The purpose of the present article is to show that neurobiological studies support suggestions from the behavioural findings that spatial learning (and, by implication, perhaps other kinds of learning too) is quasi-modular, with different brain systems operating on different kinds of sensory information. [Note that “module” does not equate to “modality” – a tactile cue and a visual cue may input into the same module (e.g., the place cell system, described below), whereas a visual stimulus like a landmark may input into different modules simultaneously (e.g., the head direction system and the beacon navigation system)]. It will be argued here that interactions within and between the modules may or may not obey the same rules as each other, making the issue of cue competition one that will ultimately have to be resolved at the level of modules – and even below that, at the level of synapses. Insofar as similar phenomena (such as, say, blocking and overshadowing) occur in different learning domains, this is not necessarily due to any universality of learning rules, but rather because these systems have independently evolved these rules as useful heuristics in those particular domains. Because of the complex different ways in which these modular systems interact, rules operating at the level of synapses or of connected structures will not necessarily translate to universal rules operating across behavioural domains, and so a universalist explanation of learning, in which given rules apply to all domains, is unlikely to be feasible. The emerging generation of learning theories needs to take the modularity of learning systems into account. Neurobiology, which provides a direct window into brain processes, has much to contribute to this endeavour.

The modularity of spatial systems as revealed by neurobiological studies

The question about whether associative learning principles are universal has been explored, by Pearce and others, in the domain of spatial processing, and this relatively well-understood system is amenable not only to behavioural but also to neurobiological analysis of the underlying learning mechanisms. Neurobiological studies conducted over many years have revealed that spatial learning tasks like the watermaze are dependent on the hippocampal system, a network that includes not only hippocampus but also input and output cortical structures such as entorhinal cortex and subiculum, together with forebrain subcortical structures including the fornix and septal nuclei. Tasks that depend on the hippocampal system seem to be those in which navigation requires flexible use of cues having constant allocentric position (position in the world) but variable egocentric position (position with respect to the animal), by virtue of the animal’s moving around. Spatial tasks are learned quickly, sometimes in only one trial (Steele & Morris, 1999). By contrast, it has emerged more recently that another class of spatial task is resistant to hippocampal damage but instead seems to depend on the striatum (Packard & McGaugh, 1996). These hippocampal-independent tasks are those in which navigation requires stereotypical responding, and in which the relevant environmental cues have a constant egocentric position at the time that behavioral decisions are made (with allocentric position being either constant *or* variable).

Some spatial tasks seem to use both systems, at different times or under different circumstances. In these tasks, navigation seems to start off in a rapidly-learned hippocampal-dependent fashion, using the allocentric position of cues but then, on overtraining, turns into a less flexible more automatic “response”-based navigation in which cues are used to guide behavioural decisions in an egocentric manner. The most famous example of this switching comes from a task introduced by Tolman, Ritchie and (Tolman,

Ritchie, & Kalish, 1946), in which animals were trained from a constant start arm on a plus-maze to always make the same body turn onto the same goal arm. If, early on in training, the animals were probed by starting them on the start arm opposite the usual one then they would make the opposite body turn so as to end up on the same goal arm as previously – guided by allocentric (world-centred) environmental cues. If the animals were overtrained, however, then on probe trials they would make the same body turn as previously, guided by egocentric (body-centred) cues and thus end up on the opposite goal arm. In an elegant set of inactivation and lesion studies, Packard and colleagues (Packard et al., 1996) have found that the “place” responding that occurs early in training depends on the hippocampus, but the “response” responding that occurs after overtraining depends on the striatum, with the remarkable consequence that hippocampal inactivation causes the animals to revert to a spatial strategy whereas striatal inactivation causes even overtrained animals to revert to a place strategy. Recent functional imaging studies in humans have confirmed that hippocampus and striatum appear to compete with each other for control of responding, with a negative correlation between hippocampal and striatal activation (Poldrack et al., 2001; Doeller, King, & Burgess, 2008).

Leaving the spatial vs. habit systems (or “modules”) aside for the moment, we now turn to a second kind of modular division within spatial processing, which has been more or less ignored by behavioural psychologists but which proves to be highly relevant to the analysis of studies such as the landmark/geometry studies reviewed by Pearce. This is the division between the map and compass systems in the brain.

The “cognitive map” theory of the hippocampus was proposed following single neuron studies revealing that hippocampal neurons are selectively active in restricted portions of the environment, with different cells preferring different regions of the environment (O’Keefe & Dostrovsky, 1971) (Figure 1A). These so-called “place cells” have been intensively studied and they have a number of interesting properties. One is that although the cells’ activity is obviously determined by environmental cues, each cell integrates information from several or many cues such that any particular cue could be removed and yet the cell would still fire in the appropriate place because the other cues in the array would collectively supply sufficient spatial information (O’Keefe & Black, 1977). It is interesting to note in passing that new cues can also be added to the array, and successfully gain control over place fields – that is, the learning of these new cues by place cells is not blocked (Barry & Muller, 2010).

The discovery of the map-like place cells prompted a search for other spatially modulated cells in nearby regions, and in the early 1990s a new class of cells was reported (Taube, Muller, & Ranck JB, 1990a) named head-direction cells (Figure 1B). These cells do not fire in a spatially restricted location but each cell does restrict its firing to a particular, very narrow range of head directions, leading to the hypothesis that these cells function like a neural compass, to inform the spatial representation about the orientation of the rat. Head direction cells have been found in a wide variety of brain regions surrounding the hippocampus including also anterior thalamus (Taube, 1995), lateral mammillary nucleus (Stackman & Taube, 1998) and retrosplenial cortex (Chen, Lin, Green, Barnes, & McNaughton, 1994). For reasons that are not yet understood, the head direction areas closely correspond to the classic Papez circuit, long known to be involved in episodic memory formation.

Because of its differing properties, the head direction system can be thought of as functionally differentiable from the place cell system and thus is to some extent a different “module”, though obviously one that interacts closely with the hippocampus. There is also a third module within the spatial mapping system, which has recently been discovered in the

medial entorhinal cortex (Fyhn, Molden, Witter, Moser, & Moser, 2004; Hafting, Fyhn, Molden, Moser, & Moser, 2005). Cells in this region fire in spatially localized patches that are arranged in grid-like fashion across the surface of the environment, hence their name “grid cells” (Figure 1C). These cells are thought to integrate linear and directional self-motion information together with contextual and landmark information, in order to anchor place cell firing to the environment (McNaughton, Battaglia, Jensen, Moser, & Moser, 2006; Jeffery, 2007b). Thus, the grid cells could be thought of as a path integration module. Of course, none of these so-called modules can be completely separated from one another as they are closely interrelated, but these regions are anatomically differentiable, which is relevant when it comes to considering how cues could compete with each other in driving these regions.

Cues that modulate neuronal processing of space – cue properties influence cue processing

Behavioural studies of cue use in spatial behaviour, of the kind discussed by Pearce (Pearce, 2009) and a number of others (e.g., (Miller & Shettleworth, 2007; Cheng et al., 2005)) have been vitally important in uncovering the mechanisms of spatial processing. However, neurobiological studies, particularly in the form of single neuron recordings made from neurons in spatial structures in the brain, have been able to add to this picture, and some of the results of these studies may help with interpretation of the complex mix of behavioural findings, particularly with respect to the associative and modular debates. One very clear finding from single neuron studies is that different cues play different roles in the spatial representation depending on characteristics such as their location, their physical characteristics (such as size) and their stability. This seems to be because different subsystems within the navigation network preferentially process different types of cue, an observation that is critical to understanding why competition effects are so multifarious in behavioural studies. The properties of cues that seem to determine how they are processed are their location, their physical characteristics and their spatial stability, as discussed below.

Location

The first cue characteristic of importance is location – specifically, whether the cue is nearby or far away. Place cells (and presumably head direction cells, though this has not yet been tested) do not respond to rotations of objects that are located very nearby, within a recording arena, but they do respond to rotations of the same objects if these are located at the periphery of the arena (Cressant, Muller, & Poucet, 1997) and head direction cells respond even better if the cues are further away (Zugaro, Berthoz, & Wiener, 2001). This division of labour between central and peripheral cues is presumably adaptive, because more distant objects provide better directional information since the rat cannot walk around them in a way that they can walk around intra-apparatus objects. Indeed, head direction cells are very responsive to landmarks that provide directional information in an otherwise unpolarised environment, and rotation of such landmarks will reliably cause head direction cells to rotate their preferred firing directions concordantly (Taube, Muller, & Ranck JB, 1990b). Taken together, these findings concerning place and head direction cells suggest that the directional component of the spatial system is predisposed to use distant cues for directional information.

By contrast with the effect of discrete landmarks on head direction cells, the boundaries of an environment seem to exert stronger effects on place cells the closer they are to an animal (Shapiro, Tanila, & Eichenbaum, 1997; Hetherington & Shapiro, 1997; Siegel, Neunuebel, & Knierim, 2008). Thus, place fields tend to be more affected by nearer walls, and fields that are located closer to walls tend to be more compact (O’Keefe & Burgess, 1996).

Physical characteristics

Distance aside, a second important characteristic of cues that influences how they are used by the spatial system is their physical nature. It was mentioned above that boundaries have different effects than discrete landmarks on the spatial system. Quite what constitutes a boundary has yet to be determined, but evidence suggests that a feature is a boundary if it somehow impedes the progress of the animal, and is also static and extended in space and time. Features that an animal can walk across unimpeded, or that are transient or focal, seem to have less influence on place cells (see (Jeffery, 2007a) for review). O'Keefe and Burgess (1996) showed that place cells typically respond to only a subset of the available walls, with different cells preferring different subsets. When boundaries of different orientations intersect in the vicinity of an animal then the environment starts to acquire a geometry, completed when the animal is completely encircled by walls that form an enclosed space. It appears that this boundedness may be important for place cells, whose firing degenerates or fragments in an unbounded space (Barry et al., 2006; Fenton et al., 2008).

While boundaries provide spatial information to place cells, another kind of physical cue type appears to provide information about context. Contextual cues may be thought of as those that define the characteristics of an environment independently of its metric qualities (see (Jeffery, Anderson, Hayman, & Chakraborty, 2004) for review), and they are characterized by their ability to induce place cells to “remap” (alter their firing patterns, as if the animal is in a completely new environment). Typical contextual cues would be the colour or smell of the environment (Anderson & Jeffery, 2003) or even perhaps more abstract properties such as the task the animal is performing there (Markus et al., 1995) or the intentions of the rat (Wood, Dudchenko, Robitsek, & Eichenbaum, 2000; Frank, Brown, & Wilson, 2000; Ferbinteanu & Shapiro, 2003). It seems that cues are contextual if they are stable and persistent, but the exact parameters that make a cue act in a contextual way have yet to be determined. As a useful interim operational definition, a cue can perhaps be regarded as contextual if its presence or absence causes place cells to remap.

Geometric and contextual cues can be differentiated on the basis of how place and other cells respond to them, a finding that is relevant to the behavioural debate about the extent to which processing of geometry depends on a “module”. Place cells certainly respond to the geometry of the environment, as evidenced by a change of activity when the environment is changed in shape (Muller & Kubie, 1987; O'Keefe et al., 1996; Lever, Wills, Cacucci, Burgess, & O'Keefe, 2002; Wills, Lever, Cacucci, Burgess, & O'Keefe, 2005). The cells also respond to contextual cues such as the colour of the enclosure, because when an environment is changed from a black box to a white one of exactly the same dimensions, the cells alter their firing patterns (Jeffery & Anderson, 2003), showing that they detected the change even though it was a non-spatial one. They also respond to changes in the smell of the enclosure, and in fact seem to process both the smell and the visual appearance in a complex way, with different cells responsive to different combinations of visual and olfactory features (Anderson et al., 2003). The population of cells as a whole thus seems to encode both geometry and context in a configural way (Anderson et al., 2006).

That said, contextual and metric cues do not seem to act in the same way. Subtle changes in geometry (i.e., those not large enough to also constitute a change in context, see below) cause the cells to shift their firing fields in a manner that is predicted by the metric properties of the change (for example, eastwards shift of a wall will induce eastwards shift of a place field, though not generally by the same amount)(O'Keefe et al., 1996). By contrast, changes in contextual cues seem to determine whether particular place fields will be expressed or not – that is, their influence is not metric so much as permissive. This distinction between metric and contextual influences on place cells is illustrated by an experiment in which the two cue types were shown to be functionally dissociable (Jeffery et

al., 2003). In this experiment, the colour of the walls and floor of a recording chamber were changed independently, and it was found, unexpectedly, that the cells altered their responsiveness to the walls (i.e., the locations of their firing fields) on the basis of the colour of the *floor*, even though the floor, being everywhere in the enclosure, does not in itself add geometric localising information. Thus, the geometric cues (from the walls, controlling where the fields are) can be functionally dissociated from the contextual cues (the floor, controlling *whether* the fields are there). As far as the geometric module hypothesis is concerned, the separation of geometric and non-geometric cues (context) indicates relatively separate processing by upstream areas within the spatial system (possibly the medial and lateral entorhinal cortices, respectively (Hargreaves, Rao, Lee, & Knierim, 2005; Hayman & Jeffery, 2008), but their combination by place cells shows that both types of cue ultimately input into the same common representation.

One final aspect of the distinction between geometric and contextual cues deserves comment, and that is that although these cues may be processed by different modules, there may also be a degree of cross-talk, inasmuch as the geometry of the environment may itself also be a context cue. When an environment is changed from, say, a square to a circle or from a rectangle to a triangle, not only do the metric properties of this environment change but so, also, do other features, such as the presence or absence of corners, the acuteness of angles, or even perhaps some computed determination of “global shape,” and these changes might act to tell the spatial system that the environment is a different one (rather than the same one slightly altered). This is an important point, because the distinction between the same environment with different cues present vs. a different environment altogether is likely to be important in any theory of learning.

As well as cues that are static features of the external world, there is a class of cue that provides important spatial information but which is internal to the animal. These are path integration cues, and comprise those spatially relevant cues (speed/distance and direction) that derive from the animal’s movement (linear and angular) through the environment. They include vestibular cues, optic (or other sensory) flow, motor efference and proprioception. Path integration cues are thought to be processed by a variety of brain systems, some as-yet-undiscovered (see (Etienne & Jeffery, 2004) and (McNaughton et al., 2006) for reviews). It is thought that the various angular and linear motion signals might be integrated by the grid cells (McNaughton et al., 2006), which can then compute a movement vector that allows place cells to update their representation.

Stability

The final characteristic of cues that affects how they are used is their stability - or, at least, their *perceived* stability. It seems that spatially stable cues can be used by both the place system and the response system, whereas unstable cues are only used by the response system (assuming they reliably predict a goal) but cannot, of course, be used to provide spatial information.

An example of the role of cue stability concerns the use by head direction cells of directional landmarks. The attachment of landmarks to head direction cells and place cells occurs quite quickly, within a few minutes of the animal’s entry into a novel environment (Taube & Burton, 1995). However, it seems that for this to happen, the landmark must seem to be stable with respect to the animal’s own internal sense of direction or it will come to be disregarded. This holds true even if (or rather, it transpires, *especially* if) the experimenters had attempted to disorient the animal prior to placing it into the novel environment. Such prior disorientation (achieved by rapidly rotating the animal) was originally assumed to make an animal more likely to use the static landmarks as directional anchors, because its own internal direction sense had been disrupted. In fact, single neuron studies reveal that

quite the reverse seems to be the case – prior disorientation makes head direction cells (and by implication, the animal's sense of direction) *less* likely to use static landmarks as cues (Knierim, Kudrimoti, & McNaughton, 1995). It seems as though the disorientation procedure realigns but does not diminish the internal direction sense, and because the landmarks thus seem to vary in orientation with respect to this sense, they are disconnected from the spatial representation in the same way that genuinely unstable landmarks are (Biegler & Morris, 1993). This is a counterintuitive finding that should prompt a reappraisal of some behavioural studies which animals were disoriented between trials (e.g., Cheng, 1986).

The above discussion has shown that cues play different roles in the spatial representation, not only because of general factors like salience or past history but also, more specifically, because they selectively influence different subsystems within the navigational network. We have identified discrete cues (landmarks), boundaries (collectively making a geometry), context and path integration. The use of some cues might be ambiguous however. For example, a large white panel on the wall of an otherwise black enclosure might be a directional landmark (if it is stable) but it might also be a contextual cue. Similarly, as discussed above, the shape of an enclosure might provide geometric information but it might also provide contextual information. A solitary landmark might be a beacon (to be used in locating a goal via a response strategy) but it might also be a directional cue, or a context cue. In the next section, we will see how not knowing what information particular cues are contributing to the spatial representation can lead to ambiguities in how to interpret the behavioural studies, which in turn can cause problems for determining the extent to which cue interactions are obeying, or failing to obey, a given set of putative learning rules.

Re-evaluating behavioural studies in light of single neuron data

How can information about neural processing of space help illuminate the behavioural data? Pearce (2009) has discussed three major issues in the spatial learning literature that have been difficult to resolve with behavioural studies. These are (1) Does spatial learning obey associative learning rules? (2) Is there a geometric module that is impervious to non-geometric (“featural”) information? and (3) Is geometric processing global or local? Single neuron data can be of great help in resolving these debates, and an attempt to do so is outlined below. Because the relevant literatures are so large, only studies that have been particularly influential, or that have posed specific conundrums, have been singled out for analysis. All three issues are discussed together, since the experiments that explore them have a high degree of overlap.

The classic test of whether a behavioural competence obeys associative learning rules is whether blocking and overshadowing can be demonstrated in this domain. Behavioural studies on blocking and overshadowing in spatial learning have produced a very confused picture, with some studies showing these effects and some studies failing to show them. These studies have been reviewed in detail by Pearce (see also (Chamizo, 2003; Cheng et al., 2005)) and they will not be re-reviewed here, but a few deserve particular scrutiny through a neurobiological lens, as it were. These are studies in which blocking and overshadowing failed to occur in the predicted manner.

The first, and most surprising, failure to find overshadowing was in the classic Cheng (1986) study of geometry vs. features, in which rats preferred to search in two geometrically equivalent corners of a rectangular enclosure for food even though the correct corner was unambiguously signalled by the presence of a unique feature (a coloured panel). This finding would not have been predicted by associative learning theory, which would suppose that the unambiguous feature, which was reliably associated with food, should come to be

strongly associated with the goal whereas the ambiguous geometry should be overshadowed. That quite the opposite was the case was suggested to imply the existence of the geometric module. By this view, geometry has intrinsically superior salience (i.e., is preferentially processed) and thus dominates even if it is less predictive of the goal.

Knowing what we now know about the neurobiology of spatial processing, how might this experiment be interpreted? There are two important features of the methodology that may be relevant for the way in which cues were processed neurally. The first is that rats were disoriented between trials, in an attempt to force them to disregard distal cues and encourage them to focus on the cues in the box. However, we saw in the Knierim *et al.* (1995) study of head direction cells, discussed earlier, that disorientation paradoxically causes intramaze landmarks to be disregarded, presumably because they conflict with the more-trusted internal sense of direction. If what is true of head direction cells is true for the animal as a whole then perhaps the disoriented rats perceived the features as varying in their location. One might therefore hypothesise that Cheng's rats were, in this particular version of the study, predisposed *not* to use the features, leaving only geometry as a reliable (albeit ambiguous) cue. A second methodological issue is that the geometric superiority was only seen in the working memory version of the task, in which the food moved from trial to trial. In the reference memory version the rats in fact learned to use the features. Here, we might suggest that with this overtraining the task switched from a place-based task to a response-based one, in which a focal stimulus became associated with an automatic, habit-like behaviour. Thus, we can suggest that Cheng may have been right that geometry is modular (can operate independently of features) but not necessarily right to also suppose an automatic, intrinsic preference for use of this module.

A second failure to find blocking is that of the Pearce *et al.* (2001) study in which the presence of an intramaze landmark not only failed to overshadow or block, but actually sometimes *potentiated* learning about the shape of the enclosure. In these experiments, rats were trained to locate hidden platforms in watermazes, in the presence or absence of one or two beacons, and then tested in the absence of any beacons to see whether they could now use the shape of the pool alone. The rats trained with beacons could do so in every case, and sometimes did even better than rats not trained with a beacon. These results were replicated in a follow-up study by Hayward *et al.* (Hayward, McGregor, Good, & Pearce, 2003), in which rats trained in a rectangular watermaze, with the platform located near a corner and with or without a nearby landmark present, were then tested in the absence of the landmark. Rats that were trained with the landmark were able to search both geometrically correct corners, showing that they must have learned about the shape of the pool, and yet during training they had learned to swim to the correct one of these corners, which was signaled by the landmark, thus showing they knew something about the landmark too. Thus, neither the landmark nor the geometry of the enclosure had overshadowed the other cue type. The control rats, trained with the platform in both corners, also searched both during testing, as would be expected.

Why might such overshadowing have failed? Based on the processes uncovered by neurobiological studies, there are several possibilities. The rats trained with the landmark may have solved it as a straightforward beacon (response) task: "swim towards the corner having the big black sphere hanging near it", but in the absence of the landmark were forced to revert to a place-based strategy. In other words, perhaps it was that response learning failed to block place learning. Alternatively, they may have solved the problem as a place task in both phases: "swim towards the north-east corner, where the orientation of the environment is indicated by the off-centre landmark". In this latter case, perhaps overshadowing failed because learning about direction does not interfere with learning about shape.

The rats trained without the landmark also, interestingly, could have formed one of two different solutions. The first is “I think the platform is in the [say] north-east corner, but because I don’t know which way is north in this symmetrical environment I’ll first guess that north is over *there*, and if that fails I’ll assume that north is the other way”. The other is “The platform is in either [say] the north-east or south-west corner: I don’t know which way round this environment is right now but I’ll assume north is that way throughout the trial, and then I will try first one corner and then the other”. In the first situation, one would predict that head direction cell activity should switch 180 degrees as the rat switches its search (for what it thinks is a single goal location) from the first corner to the second. In the second situation one would predict that the head direction cells would stay constantly oriented, with the rat trying what it thinks are two different possible goal locations. What actually happens in this situation is a matter for experimental testing. The point is that the same behavioural outcome may derive from the operation of fundamentally different strategies at a neural level, and to understand the associative processes one needs to know which strategy is actually operative, perhaps by recording neural activity during execution of these tasks.

A second experiment in this study also found a failure of a landmark to block geometric learning. Rats were trained in the rectangular pool with the landmark and platform randomly located in any one of the four corners, and then trained again in a triangular pool with the landmark and platform in a constant corner. The prediction was that because rats had learned in the rectangle that the landmark predicted the platform and the shape did not, blocking should occur such that in the new environment, the animals would not learn about shape because the goal was already perfectly predicted by the landmark. In fact, this did not happen: after training in the new pool, when the landmark was removed the rats readily searched in the correct corner (indicated by the geometry).

Again, what neurobiological processes might underlie the learning of these tasks? In the rectangular pool, one might hypothesise that the animals had acquired a response (beacon-based, probably striatal) strategy for locating the platform. But, why did this strategy fail to block place learning in the triangular pool? There are two possibilities. The first, a possibility that also arose above, is that perhaps response processing simply does not block place processing. The second possibility is that perhaps blocking between landmarks and shape *can* occur, but did not in this experiment because the shape change also constituted a change in context. We know from place cell studies that change in shape of an environment is a context cue sufficiently salient to induce remapping (Lever et al., 2002), so such changes are likely to be influential in spatial processing. Associative rules like the Rescorla-Wagner rule (Rescorla & Wagner, 1972) predict that a salient change in context should cause a degree of unblocking, because the context cues in the first environment become somewhat associated with the goal and thus occupy some of the available associative strength. The change in context (in the new training enclosure) would, as it were, release this associability so that new connections (between the new shape and the goal) could form. This possibility is not only feasible but actually rescues the associative learning account of behaviour. Again, however, the message is that without knowing the underlying neural representation of a problem, it is hard to draw conclusions about the relevant associative processes merely from behavioural observations.

In a further set of studies, Pearce and colleagues (Pearce, Graham, Good, Jones, & McGregor, 2006; Graham, Good, McGregor, & Pearce, 2006) explored the interaction between shape and non-shape cues (features) using, instead of landmarks, coloured walls to provide additional information in the environment. Rats were able to use either the location of the coloured walls, the geometry, or (sometimes) both to locate the hidden platform. Interestingly, it was found that rats trained with colour and shape together were able to use

either cue type when deprived of the other, showing that shape had not blocked colour and vice versa. Furthermore, they also did better than the rats trained with only shape being relevant, showing a positive interaction between shape and colour. This unexpected potentiation of landmarks by colour was explained by Pearce as indicative of between-cue associations (Pearce, 2009). However, it is also reminiscent of the positive interaction between internal directional orientation and directional landmarks in the Knierim et al. (1995) study discussed earlier. Such findings suggest that the directional subsystem learns about cues more easily when they support each other, perhaps because animals having stable directional (e.g., colour in this case) cues are more quickly able to orient when placed in the environment, and thus more easily able to learn the goal location. In this light, it is worth noting that rats have very poor, monochrome vision and that the distinction between an acute and an obtuse angle may be very subtle and difficult for them, whereas the distinction between black and white walls is likely to be highly salient as brightness is readily, and in fact almost preferentially processed (e.g., see (Minini & Jeffery, 2006)). This could explain why rats given shape alone as a cue showed much slower learning and barely-above-chance performance at asymptote (Graham et al. 2006, Figure 1). Pursuing their line of enquiry further, Pearce et al (2006) then trained rats in either a kite-shaped or rectangular watermaze, with either all-white walls or a mixture of black and white walls. Again, they found that under some circumstances cues blocked each other but in other circumstances they did not, confirming that there is no simple overarching principle by which cue interaction effects can be predicted.

A final failure to find blocking between landmarks and geometry was a human study, conducted by Doeller and Burgess (2008). Subjects in a virtual reality task were required to find objects whose position could be encoded relative to a landmark, to the boundary of the environment, or both. When these cues were presented in a blocking paradigm and then tested in isolation, it was found that landmarks blocked each other, but did not block the boundary.

What are we to make of this profusion of behavioural findings, with some studies (and only a small subset of them were discussed here) finding blocking and some finding potentiation? (Wall, Botly, Black, & Shettleworth, 2004) have drawn a succinct and plausible conclusion, echoed here, which is that “[spatial] subsystems include dead reckoning, beacon (or cue) learning, landmark learning, response or habit learning, and learning the overall shape of the local space formed by enclosing surfaces”. This conclusion derived from their analysis of both behavioural and neurobiological literature, and seems to lucidly capture the essence of the problem: there actually are a number of interacting subsystems, and one needs to know which ones are operative before being able to understand how cues are being used, and competing (or not).

A strategy for interpreting behavioural studies

In the spatial domain, single neuron studies have led to the derivation of two overarching conclusions. The first is that the navigation system is quasi-modular, inasmuch as there are different subsystems that contribute different processing roles – one subsystem computes a directional signal, for example, while another one processes context and tries to decide whether the animal is in a new context or an altered familiar one. The second conclusion is that all cues are not equal – different cues play different roles, depending on their sensory modalities, location and sensory qualities, and some cues may play more than one role. Thus, to untangle the web of conflicting behavioural studies, one needs to know not only what cues the animal is using but also what function(s) each cue is contributing to a particular task.

Moving on from these two general conclusions, a number of specific conclusions have arisen from neurobiological studies of spatial cognition. They could be thought of as signposting potential points of ambiguity in a behavioural experiment, and they are as follows:

(1) Learning to navigate to a goal involves co-operation and/or competition between (at least) two overarching systems, the hippocampal place system and the striatal response system. To interpret a particular experiment one needs to know which system the animal is using, at a given moment, and also to recognize that this may change over time, or if the experimenter replaces one cue by another.

(2) Focal cues (“landmarks”) sometimes act as beacons, particularly if they are reliably located at or near a goal and the animal has been trained for some time. They may alternatively act as directional cues, if they are located further away, and seem to be stable (i.e., if they do not move and if the animal was not disoriented when learning about them). Since beacon use probably is a striatal response function while use of a directional cue is more likely a head-direction/hippocampal one, it is necessary to consider what role a particular landmark plays in a given setting. It also needs to be considered that rodent eyesight is poor and that two cues that look very different to a primate experimenter may be indistinguishable to a rodent subject, particularly if they are far away.

(3) Geometric cues (i.e., those cues contributed by environmental objects of fixed dimension and fixed relative location) can act in a variety of ways. They may act in the way that Cheng (1986) originally proposed, which is, by virtue of their global shape, to indicate the location of a goal which the animal has to extract by metrically processing the cues. They may alternatively, as Pearce (2009) has argued persuasively, act as beacons if the animal uses local properties in a snapshot-like way (e.g., long wall to the left, short wall to the right) (see also (Cheng, 2008a)). They may also, and this is rarely considered by behavioural experimenters, act as contextual cues. For example, if a rectangular enclosure is changed to a square, which is sometimes done to remove geometric localizing information in a probe trial, the animal may consider this to be a change in context. Furthermore, this perceived context change may be sufficiently great as to affect the animal’s use of previously learned information. To put this loosely, it may be that the animal thinks it is in a different environment and is doubtful about whether the things it learned in the previous environment are still relevant. To put it in associative learning terms, training in one context may, by the predictions of the Rescorla-Wagner rule (Rescorla & Wagner, 1972), stop cues from gaining full associative strength because the context sequesters some of the associative strength. When this context is replaced, the remaining cues are at submaximal associative strength and there is room left for a new cue to gain some associativity (i.e., there will be a degree of unblocking). Thus, an experimenter needs to consider whether the environmental change in a probe trial is sufficiently large as to constitute a context change – in which case all bets are off as to how the animal will use the cues in this new context.

One more role of geometric cues that is often overlooked is their contribution to the head direction system. Little is known about whether head direction cells can process geometry, but preliminary indications suggest that they can (Taube et al., 1990b; Hartley, Trinkler, & Burgess, 2004). This of course may, again, be either via global processing or by a local snapshot-type processing. Thus, when a change is made to the shape of an enclosure in a curtained-off environment, the head direction cells may receive different information (or even no information) about direction, or the information they were extracting from the geometry might now conflict with the information from a landmark that they were also using.

(4) Path integration is also a cue, and appears to be a highly trusted one, and furthermore it seems to be constantly operative, certainly in the angular domain. Single neuron studies of head direction cells tell us that the cells are always active and always consistent with each other – that is, the head direction system seems to maintain a continual representation of which way round an animal is facing, even if this representation can be altered (not removed) by disorientation. There is no evidence that an experimenter can switch the path integrator off, as it were, by (for example) disorienting an animal. This needs to be remembered when manipulating landmarks: path integration is still present, and if the landmarks conflict with the path integrator too often then the animal may just learn to disregard them altogether.

One can see that given this mish-mash of interacting co-operating/competing systems, all trying to operate on the same or different cues, it becomes difficult to falsify any particular theoretical position using a behavioural approach alone, because an investigator can always appeal to some ambiguity in how a cue is being used that might account for discrepant results. Studies of single neurons by themselves are insufficient too, because the mere fact that a set of neurons responds to a change in the environment does not mean that the animal is using those neurons to solve the task. The bottom line is that a co-operation between both approaches is needed. Behavioural and physiological researchers collectively need to rescind their hitherto isolated approaches, and collaborate to produce an orchestrated attack on the problem of how animals learn.

Conclusion – Universalist vs quasi-modular accounts of learning

Classical associative learning theory evolved to account for learning in a particular domain – predictive learning of temporally sequential events (such as stimuli and responses, or stimuli and other stimuli). The question at hand is whether these principles can be applied to all learning, or whether they are domain-specific. Evidence from behavioural and neurobiological studies points to the quasi-modularity of learning systems (“quasi” because the so-called modules are not independent), in which different kinds of learning recruit different sets of brain systems.

Why does quasi-modularity challenge the universality position? The reason is that universalist learning rules are stimulus- or outcome-independent, which is to say that they can be described in abstract terms (A or B, X or Y etc) in which the exact nature of the stimulus is irrelevant – it does not matter whether A is a tone, a light or a food flavour, the rule is the same. Neurobiology tells us that stimuli cannot be interchanged in this way – a directional landmark is not necessarily equivalent to a coloured panel or a wall, and the exact nature of the stimulus affects which module it feeds into, and therefore what path it takes through the, as it were, algorithmic network, as in the example in Figure 2. Thus, an associative rule describing how a landmark affects a rat’s calculation of food location must take into account whether that landmark is setting the head direction system, or acting as a beacon, or perhaps both. To fully explain behaviour in “rule” terms we need complex networks of rules that take into account the modularity of the brain. It may well turn out that rules operating between stimuli within a module are the same as, or at least similar to, those describing cue interactions in a different module, and so there may be some generality in this regard. This will need to be determined experimentally. We will know that we have a properly explanatory model when we can not only describe how an animal’s behaviour arose *post hoc* but can also predict it.

Knowing what we now know about the specialised neural organisation that supports spatial learning, it is time to set aside global views such as “all learning is error-correcting” or “all self-localisation is geometric.” Rather, we need to look into the brain and determine which

areas contribute which function in the generation of a given behavior, an enterprise that will require a collaboration between behavioural scientists and neuroscientists. The end result will be a new generation of associative learning model in which descriptions of behavior will take into account not only the physical nature of the stimuli that drive that behaviour, but also the structure of the network in which they are processed and their functional role(s) in the network.

Acknowledgments

The work is supported by grants from the Wellcome Trust (GR083540AIA) and the European Commission Framework 7 ("Spacebrain"). The author would like to thank Bernard Balleine, Robin Murphy, Marko Nardini and Clare Press for helpful comments on the manuscript.

Grant sponsors: Wellcome Trust (GR083540AIA), EU FP7 ('Spacebrain')

References

- Anderson MI, Jeffery KJ. Heterogeneous modulation of place cell firing by changes in context. *J. Neurosci.* 2003; 23:8827–8835. [PubMed: 14523083]
- Anderson M, Killing S, Morris C, O'Donoghue A, Onyiaha D, Stevenson R, et al. Behavioral correlates of the population coding of spatial context. *Hippocampus.* 2006; 16:730–742. [PubMed: 16921500]
- Barry C, Lever C, Hayman R, Hartley T, Burton S, O'Keefe J, et al. The boundary vector cell model of place cell firing and spatial memory. *Reviews in the Neurosciences.* 2006; 17:71–97. [PubMed: 16703944]
- Barry J, Muller R. Updating the hippocampal representation of space: Place cell firing fields are controlled by a novel spatial stimulus. *Hippocampus.* 2010 in press.
- Biegler R, Morris RG. Landmark stability is a prerequisite for spatial but not discrimination learning. *Nature.* 1993; 361:631–633. [PubMed: 8437622]
- Chamizo VD. Acquisition of knowledge about spatial location: assessing the generality of the mechanism of learning. *Q J Exp Psychol B.* 2003; 56:102–113. [PubMed: 12623541]
- Chen LL, Lin LH, Green EJ, Barnes CA, McNaughton BL. Head-direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. *Exp. Brain Res.* 1994; 101:8–23. [PubMed: 7843305]
- Cheng K. A purely geometric module in the rat's spatial representation. *Cognition.* 1986; 23:149–178. [PubMed: 3742991]
- Cheng K. Whither geometry? Troubles of the geometric module. *Trends Cogn Sci.* 2008; 12:355–361. [PubMed: 18684662]
- Cheng K, Newcombe NS. Is there a geometric module for spatial orientation? Squaring theory and evidence. *Psychon. Bull. Rev.* 2005; 12:1–23. [PubMed: 15945200]
- Cressant A, Muller RU, Poucet B. Failure of centrally placed objects to control the firing fields of hippocampal place cells. *J Neurosci.* 1997; 17:2531–2542. [PubMed: 9065513]
- Doeller CF, Burgess N. Distinct error-correcting and incidental learning of location relative to landmarks and boundaries. *Proc Natl Acad Sci U S A.* 2008; 105:5909–14. [PubMed: 18413609]
- Doeller CF, King JA, Burgess N. Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proc. Natl. Acad. Sci U.S.A.* 2008; 105:5915–5920. [PubMed: 18408152]
- Etienne AS, Jeffery KJ. Path integration in mammals. *Hippocampus.* 2004; 14:180–192. [PubMed: 15098724]
- Fenton AA, Kao HY, Neymotin SA, Olypher A, Vayntrub Y, Lytton WW, et al. Unmasking the CA1 ensemble place code by exposures to small and large environments: more place cells and multiple, irregularly arranged, and expanded place fields in the larger space. *J Neurosci.* 2008; 28:11250–11262. [PubMed: 18971467]

- Ferbinteanu J, Shapiro ML. Prospective and retrospective memory coding in the hippocampus. *Neuron*. 2003; 40:1227–1239. [PubMed: 14687555]
- Frank LM, Brown EN, Wilson M. Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron*. 2000; 27:169–178. [PubMed: 10939340]
- Fyhn M, Molden S, Witter MP, Moser EI, Moser MB. Spatial representation in the entorhinal cortex. *Science*. 2004; 305:1258–1264. [PubMed: 15333832]
- Golob EJ, Taube JS. Head direction cells in rats with hippocampal or overlying neocortical lesions: evidence for impaired angular path integration. *J Neurosci*. 1999; 19:7198–7211. [PubMed: 10436073]
- Graham M, Good MA, McGregor A, Pearce JM. Spatial learning based on the shape of the environment is influenced by properties of the objects forming the shape. *J Exp Psychol Anim Behav Process*. 2006; 32:44–59. [PubMed: 16435964]
- Hafting T, Fyhn M, Molden S, Moser MB, Moser EI. Microstructure of a spatial map in the entorhinal cortex. *Nature*. 2005; 436:801–806. [PubMed: 15965463]
- Hargreaves EL, Rao G, Lee I, Knierim JJ. Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. *Science*. 2005; 308:1792–1794. [PubMed: 15961670]
- Hartley T, Trinkler I, Burgess N. Geometric determinants of human spatial memory. *Cognition*. 2004; 94:39–75. [PubMed: 15302327]
- Hayman M, Jeffery KJ. How heterogeneous place cell responding arises from homogeneous grids-A contextual gating hypothesis. *Hippocampus*. 2008; 18:1301–1313. [PubMed: 19021264]
- Hayward A, McGregor A, Good MA, Pearce JM. Absence of overshadowing and blocking between landmarks and the geometric cues provided by the shape of a test arena. *Q J Exp Psychol B*. 2003; 56:114–126. [PubMed: 12623542]
- Hetherington PA, Shapiro ML. Hippocampal place fields are altered by the removal of single visual cues in a distance-dependent manner. *Behav. Neurosci*. 1997; 111:20–34. [PubMed: 9109621]
- Jeffery KJ. Integration of the sensory inputs to place cells: what, where, why, and how? *Hippocampus*. 2007a; 17:775–785. [PubMed: 17615579]
- Jeffery KJ. Self-localization and the entorhinal-hippocampal system. *Curr. Opin. Neurobiol*. 2007b; 17:684–691. [PubMed: 18249109]
- Jeffery KJ, Anderson MI. Dissociation of the geometric and contextual influences on place cells. *Hippocampus*. 2003; 13:868–872. [PubMed: 14620882]
- Jeffery KJ, Anderson MI, Hayman R, Chakraborty S. A proposed architecture for the neural representation of spatial context. *Neurosci Biobehav Rev*. 2004; 28:201–218. [PubMed: 15172764]
- Knierim JJ, Kudrimoti HS, McNaughton BL. Place cells, head direction cells, and the learning of landmark stability. *J. Neurosci*. 1995; 15:1648–1659. [PubMed: 7891125]
- Le Pelley ME. The role of associative history in models of associative learning: a selective review and a hybrid model. *Q J Exp Psychol B*. 2004; 57:193–243. [PubMed: 15204108]
- Lever C, Wills T, Cacucci F, Burgess N, O'Keefe J. Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature*. 2002; 416:90–94. [PubMed: 11882899]
- Markus EJ, Qin YL, Leonard B, Skaggs WE, McNaughton BL, Barnes CA. Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J. Neurosci*. 1995; 15:7079–7094. [PubMed: 7472463]
- McNaughton BL, Battaglia FP, Jensen O, Moser EI, Moser MB. Path integration and the neural basis of the 'cognitive map'. *Nat. Rev. Neurosci*. 2006; 7:663–678. [PubMed: 16858394]
- Miller NY, Shettleworth SJ. Learning about environmental geometry: an associative model. *J Exp Psychol Anim Behav Process*. 2007; 33:191–212. [PubMed: 17620021]
- Minini L, Jeffery KJ. Do rats use shape to solve "shape discriminations"? *Learning and Memory*. 2006; 13:287–297. [PubMed: 16705141]
- Muller RU, Kubie JL. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci*. 1987; 7:1951–1968. [PubMed: 3612226]
- O'Keefe J, Black AH. Single unit and lesion experiments on the sensory inputs to the hippocampal cognitive map. *Ciba. Found. Symp*. 1977:179–198.

- O'Keefe J, Burgess N. Geometric determinants of the place fields of hippocampal neurons. *Nature*. 1996; 381:425–428. [PubMed: 8632799]
- O'Keefe J, Dostrovsky J. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res*. 1971; 34:171–175. [PubMed: 5124915]
- O'Keefe, J.; Nadel, L. *The Hippocampus as a Cognitive Map*. Clarendon Press; Oxford: 1978.
- Packard MG, McGaugh JL. Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiol.Learn.Mem*. 1996; 65:65–72. [PubMed: 8673408]
- Pearce JM, Ward-Robinson J, Good M, Fussell C, Aydin A. Influence of a beacon on spatial learning based on the shape of the test environment. *J.Exp.Psychol.Anim Behav.Process*. 2001; 27:329–344. [PubMed: 11676084]
- Pearce JM. The 36th Sir Frederick Bartlett Lecture: An associative analysis of spatial learning. *Q J Exp Psychol (Colchester)*. 2009; 62:1665–84. [PubMed: 19418377]
- Pearce JM, Graham M, Good MA, Jones PM, McGregor A. Potentiation, overshadowing, and blocking of spatial learning based on the shape of the environment. *J Exp Psychol Anim Behav Process*. 2006; 32:201–214. [PubMed: 16834489]
- Poldrack RA, Clark J, Pare-Blagoev EJ, Shohamy D, Creso MJ, Myers C, et al. Interactive memory systems in the human brain. *Nature*. 2001; 414:546–550. [PubMed: 11734855]
- Rescorla, RA.; Wagner, AR. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black, AH.; Prokasy, WF., editors. *Classical Conditioning II: Current Research and Theory*. Appleton-Century-Crofts; New York: 1972. p. 64-99.
- Shapiro ML, Tanila H, Eichenbaum H. Cues that hippocampal place cells encode: dynamic and hierarchical representation of local and distal stimuli. *Hippocampus*. 1997; 7:624–642. [PubMed: 9443059]
- Siegel JJ, Neunuebel JP, Knierim JJ. Dominance of the proximal coordinate frame in determining the locations of hippocampal place cell activity during navigation. *J Neurophysiol*. 2008
- Stackman RW, Taube JS. Firing properties of rat lateral mammillary single units: head direction, head pitch, and angular head velocity. *J.Neurosci*. 1998; 18:9020–9037. [PubMed: 9787007]
- Steele RJ, Morris RG. Delay-dependent impairment of a matching-to-place task with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5. *Hippocampus*. 1999; 9:118–136. [PubMed: 10226773]
- Taube JS. Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *J.Neurosci*. 1995; 15:70–86. [PubMed: 7823153]
- Taube JS, Burton HL. Head direction cell activity monitored in a novel environment and during a cue conflict situation. *J.Neurophysiol*. 1995; 74:1953–1971. [PubMed: 8592189]
- Taube JS, Muller RU, Ranck JB J. Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J.Neurosci*. 1990a; 10:420–435. [PubMed: 2303851]
- Taube JS, Muller RU, Ranck JB J. Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *J.Neurosci*. 1990b; 10:436–447. [PubMed: 2303852]
- Tolman EC. Cognitive maps in rats and men. *Psychol.Rev*. 1948; 40:40–60.
- Tolman EC, Ritchie BF, Kalish D. Studies in spatial learning: II. Place learning versus response learning. *Journal of Experimental Psychology*. 1946; 37:385–392.
- Wall PL, Botly LC, Black CK, Shettleworth SJ. The geometric module in the rat: independence of shape and feature learning in a food finding task. *Learn.Behav*. 2004; 32:289–298. [PubMed: 15672824]
- Wills TJ, Lever C, Cacucci F, Burgess N, O'Keefe J. Attractor dynamics in the hippocampal representation of the local environment. *Science*. 2005; 308:873–876. [PubMed: 15879220]
- Wood ER, Dudchenko PA, Robitsek RJ, Eichenbaum H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron*. 2000; 27:623–633. [PubMed: 11055443]

Zugaro MB, Berthoz A, Wiener SI. Background, but not foreground, spatial cues are taken as references for head direction responses by rat anterodorsal thalamus neurons. *J Neurosci.* 2001; 21:RC154. [PubMed: 11425881]




Figure 1.

(A) The firing pattern of a typical hippocampal place cell, recorded as a rat foraged in an environment 70 cm square with 50 cm high walls. The path of the rat is shown by the stippling and the action potentials of the cell are shown as small squares, each plotted at the location of the rat at the time it occurred. Note that this particular cell prefers to fire slightly south-east of the centre of the box, and was silent elsewhere. (B) The firing pattern of a typical head direction cell (adapted from (Golob & Taube, 1999)). This cell fired everywhere in the environment (not shown), but only when the rat's head was pointing in a particular, highly restricted direction. (C) The firing pattern of a typical entorhinal grid cell (adapted from Hafting et al., 2005), shown in the same format as the place cell in (A). Note that like the place cell, this grid cell fired in a spatially localized way, but unlike the place cell it had multiple, regularly-spaced firing locations.

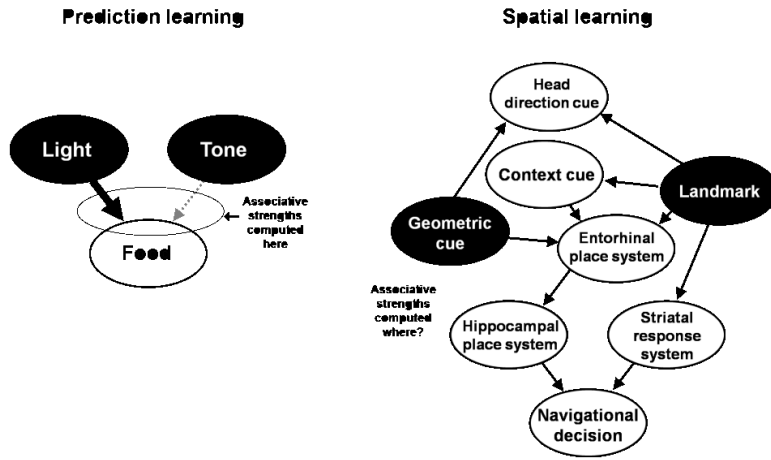


Figure 2. Why the quasi-modularity of learning systems poses challenges to universal rules of description. Left: A schematic of cue interaction in prediction learning, where the associative strength of a stimulus depends on the summed associative strength of all stimuli impinging on the output. In this example, developed from Figure 1, learning about the light has been blocked by prior learning about the food, and the calculation of associative strengths occurs at the place where light and tone processing pathways converge. Right: Spatial learning involves a number of sub-modules. At which place in this network could an equivalent summed-associative-strength calculation occur? The complexity of most brain systems suggests that a universal-rules description of learning processes is unlikely to capture behaviour in a usefully predictive manner.