

What Do We Know about Spatial Navigation, and What Else Could Model-Based fMRI Tell Us?

Adam Tyson, MSc

Faculty of Mathematical and Physical Sciences, University College London, Gower Street, London, United Kingdom.

Spatial navigation, or the ability to remember and navigate environments, is an important skill for humans and animals. It has inspired a great deal of research, including neuroimaging studies of humans and single-unit recordings of animals. Recent advances in computational modeling have enabled spatial navigation in humans and animals to be investigated in a more precise and detailed manner. More specifically, computational models allow us to estimate theoretical parameters associated with spatial navigation, and model-based fMRI can be used to investigate the neural correlates of these parameters.

This review addresses the literature on spatial navigation beginning with reviewing lesion and animal studies of spatial cognition. Imaging studies of spatial memory and navigation in humans, including structural imaging, and more-complex functional imaging studies involving virtual reality are then discussed. Particular emphasis is placed on computational studies of behavior involving reinforcement learning models and model-based fMRI. Finally, the advantages of model-based fMRI for investigating the neural basis of spatial navigation in humans are discussed.

INTRODUCTION

At one time or another we have each become lost—maybe in a new city, heading in the wrong direction or walking in circles on the way to the hotel. In contrast, most of us can travel to and from work each day without any problems, often arriving with little recollection of the journey we took and the decisions we made along the way. Remembering and navigating environments is of great importance for humans and animals alike, yet we often take it for granted. We tend not to appreciate our ability to navigate environments until we get lost in a new city, or when our ability to navigate is compromised by Alzheimer's disease (Henderson, Mack, & Williams, 1989) or other forms of dementia.

In this review I will discuss advances in the study of spatial navigation, including results from experiments in both animals and humans. Many methods have been used, including behavioral, neuropsychological, electrophysiological, neuroimaging, and computational modeling. I will introduce reinforcement learning, which is one aspect of theoretical neuroscience that has only recently been applied in studies of spatial navigation in humans and animals.

Navigation has been studied for a long time; it had many early breakthroughs, such as those of Tolman (1948), who interpreted both his own results (Tolman, Ritchie, & Kalish, 1946) and those of others (e.g., Blodgett, 1929) as evidence that a rat has an internal allocentric representation of space, or a cognitive map of its environment. Many of the major breakthroughs in our understanding of the neural representation of space—notably the discovery of place cells—have come from animals. These pyramidal cells in the rat hippocampus fire selectively in particular areas of the animal's environment (O'Keefe & Dostrovsky, 1971) and have been interpreted as a possible neural basis for Tolman's cognitive map (O'Keefe & Nadel, 1978), allowing

the animal to navigate around obstacles or take shortcuts. More recently, head-direction cells, first found in the post-subiculum (Taube, Muller, & Ranck, 1990), and entorhinal grid cells (Hafting, Fyhn, Molden, Moser, & Moser, 2005), were discovered, the latter of which may form the basis of a path integration-based representation of the animal's environment. It is not completely clear how these findings relate to human navigation, although work has been done to find evidence for homologues of these cells in humans—for example, place cells (Ekstrom et al., 2003) and grid cells (Doeller, Barry, & Burgess, 2010).

SPATIAL MEMORY AND NAVIGATION

Neural Substrates of Spatial Memory

Since the case of patient H.M., who underwent a bilateral medial temporal lobectomy for intractable epilepsy, the medial temporal lobe (MTL), and the hippocampus in particular, have been associated with episodic memory (Scoville & Milner, 1957). In concordance with the discovery of place cells, and the idea of a cognitive map, the hippocampus is thought to be involved in spatial memory in animals (O'Keefe & Nadel, 1978; Morris, Garrud, Rawlins, & O'Keefe, 1982) and humans (Maguire, Burke, Phillips, & Staunton, 1996; Spiers, Burgess, Hartley, Vargha-Khadem, & O'Keefe, 2001). Hippocampal lesions have been shown to cause deficits in spatial memory in rodents (Morris et al., 1982), and electrical stimulation of the entorhinal cortex, an MTL structure associated with the hippocampus, and its main interface with the neocortex, has been linked to improved spatial memory in mice, possibly related to improved adult neurogenesis in the dentate gyrus (Stone et al., 2011).

In humans, many different methods have been employed to discover the exact relation of the hippocampus and other brain structures with spatial memory and navigation.

A voxel-based morphometry study showed structural differences between licensed London taxi drivers, who must pass a rigorous test of their knowledge of London roads (and thus are expected to have better spatial memory), and controls (Maguire et al., 2000). It was found that the taxi drivers had larger posterior hippocampi and smaller anterior hippocampi compared with controls. This change in hippocampal size was also found to correlate linearly with the number of years of experience the taxi driver had (Maguire, Woollett, & Spiers, 2006). In concordance with rodent-study results, electrical stimulation of the entorhinal cortex in epileptic patients undergoing invasive recordings prior to surgery resulted in quicker and more-accurate navigation in a simulated environment (Suthana et al., 2012). It is tempting therefore perhaps to link these two studies with the rodent study connecting adult neurogenesis with improved spatial memory to suggest that increased adult neurogenesis in the dentate gyrus results in a greater number of hippocampal cells, underpinning the remarkable talents of the London taxi drivers. This possible mechanism for acquisition of spatial memories has much support; adult-generated dentate gyrus cells are preferentially recruited into neural networks associated with spatial memories (Kee, Teixeira, Wang, & Frankland, 2007), and many models have been proposed to link adult neurogenesis with hippocampal learning (e.g., Becker, 2005; Aimone, Wiles, & Gage, 2006). However appealing this theory may be, there has been no conclusive supporting evidence for it.

None of the studies mentioned provides any insight into how humans or animals use these structures to know where they are and how to navigate to a goal location. Place cells are clearly an important element of spatial cognition and navigation, but there is a limit to what we can find out using experimental animals. Single-unit recordings in freely moving animals produce sensory-motor confounds; these can be controlled better by using human subjects, who are also assumed to be better at navigating and making decisions within their environments. However, one issue when studying navigation is the scale of the problem; naturally, humans navigate in large-scale environments, something that is difficult to reproduce in a controlled laboratory setting. Traditional tabletop tests of spatial memory do not accurately test natural navigation (Maguire et al., 1996), as the subjects are required to solve the problems from different viewpoints or in different reference frames from the ones they would naturally employ. Natural navigation tasks are more realistic, but they present problems when the researchers are trying to control between subjects, or accurately record performance spatially and temporally. One solution is to use virtual reality (VR).

Functional Imaging of Navigation

One of the major advantages of VR is the possibility of combining it with other techniques; its nature allows the subject to explore a virtual environment on a screen, while remaining still enough to allow functional images or single-cell recordings to be taken. Some early VR studies of spatial memory (e.g., Aguirre, Detre, Alsop, & D'Esposito,

1996) showed activation in certain brain areas (parahippocampus and associated cortex). However, it is difficult to break down the activation patterns of these studies to find the particular activity underpinning the task. Further studies have illuminated the function of different areas during navigation, using, for example, positron emission tomography (PET) while participants navigated a complex but previously experienced VR town (Maguire et al., 1998). The subjects underwent four different navigation tasks, allowing the authors to find that the participants' speed moving through the environment was associated with caudate activation, activity in the right hippocampus was associated with navigation accuracy, and activity in the left hippocampus was associated with navigation success. Bilateral medial and right inferior parietal activation corresponded with movement through the environment, and prefrontal activation was associated with success in navigating around blocked routes. However, due to the technique used (PET), between-subject effects could not be distinguished.

Recently, the use of PET has decreased in favor of functional magnetic resonance imaging (fMRI), which has many advantages, such as higher temporal resolution and very high spatial resolution. Using fMRI, Hartley, Maguire, Spiers, and Burgess (2003) expanded on the PET experiments, finding that in successful navigators, anterior hippocampal activation was correlated with way-finding, and caudate activation was correlated with route following (hence the correlation with speed found in Maguire et al., 1998).

As well as finding which brain areas are active during navigation, more-recent experiments have sought to determine which aspects of navigating in a virtual environment correspond to the detected activity. Correlations between hippocampal activity and navigation relying on spatial memory of the environment, and between parahippocampal activity and navigation relying on contextual memory (that is, the relationships between landmarks—"the post office is to the left of the statue"), have been found using variants of a learned environment (Rauchs et al., 2008). Functional segregation of the MTL at different phases of navigation has been investigated by testing subjects in variants of a learned VR environment (Xu, Evensmoen, Lehn, Pintzka, & Häberg, 2010). The authors found that anterior MTL (anterior hippocampus, entorhinal cortex, and anterior parahippocampal cortex) was active only during the initial phase of navigation, involving self-localization and planning routes (as reported by the participants), and the posterior MTL (posterior hippocampal and posterior parahippocampal cortex) was active throughout navigation, presumably corresponding to processing spatial information relating to the subjects' current position within the environment.

Using a similar but much more detailed method than Xu et al. (2010), Spiers and Maguire (2006) sought to investigate the neural activity corresponding to more-detailed aspects of navigation as their subjects (taxi drivers) drove around London in response to requests from customers. After the subjects finished the task and left the scanner, they were

immediately shown a replay of their navigation, and were interviewed to discover what they were thinking at different stages of the navigation. The verbal report protocol used with the subjects following the scan allowed the authors to break down the task into many more subcomponents than previous studies had, including visual inspection, action planning, and simply coasting. The authors found that, during the initial planning of the route, there was activation in the whole spatial-navigation network, including the hippocampus, as well as activation in lateral and medial prefrontal areas. When subjects altered their route during the journey, activation was seen in retrosplenial and right parietal cortices as well as prefrontal areas. The subjects interviewed reported expecting particular routes or landmarks; when these expectations were fulfilled, the retrosplenial and posterior parietal cortices were active. However, if, for example, they encountered a blocked route, the right lateral prefrontal cortex became active, supporting previous studies linking this area to detecting violations of expectations (e.g., Corlett et al., 2004).

These studies provide an insight into how the brain keeps track of our position as we move through space, but we do not know as much about how we navigate toward a goal location. Studies such as Spiers and Maguire (2007) provide evidence for internal metrics of goal location and distance, supporting models (e.g., Burgess, Jackson, Hartley, & O'Keefe, 2000) of how organisms navigate to a goal. This does not, however, provide any evidence for how the brain makes decisions during navigation, particularly at vital points such as when the route is blocked. Computational models have been developed to explain how organisms make decisions, and by combining these with functional imaging techniques, it is possible to discover how decision processes are carried out in different regions of the brain during navigation. These models, and their implementation along with functional imaging, will be discussed at greater length below.

As sophisticated as fMRI techniques have become, they are particularly limited by their temporal resolution. As such, more-direct measures of neural activity, such as electroencephalography (EEG) and magnetoencephalography (MEG), have been used to study the association between navigation and high-frequency brain activity, such as the theta rhythm. The theta rhythm has been linked to spatial behavior in rodents, and there is evidence from EEG (Kahana, Sekuler, Caplan, Kirschen, & Madsen, 1999) and MEG (de Araújo, Baffa, & Wakai, 2002) that these theta oscillations are linked to navigation in humans as well as lower mammals. More recently, MEG has been used to determine the function of these theta oscillations in human navigation. Cornwell, Johnson, Holroyd, Carver, and Grillon (2008) used a virtual Morris water maze and found that anterior hippocampal theta was implicated in the encoding of the spatial environment, and posterior hippocampal theta was highly correlated with navigation performance.

There is now good evidence for the neural basis of Tolman's

(1948) cognitive map, with location-specific hippocampal cells observed, and structural changes in the human hippocampus that correlate with spatial-memory abilities. However, because of the very nature of the problem, navigation in humans is hard to test, and this has led to the development of VR environments to investigate navigation in a controlled manner. VR has been successfully combined with functional imaging, and the hippocampus has been consistently linked with navigation accuracy, particularly in the early stages of navigation, in which recall of memory is most vital. Electromagnetic methods (EEG and MEG) have also been used, and a link made between the hippocampal theta rhythm and navigation. These results show that the hippocampus is almost certainly responsible for spatial cognition in animals and humans, but there is still much to be discovered. Less is known about how we make decisions during navigation, but this is where computational models may help us answer these questions.

COMPUTATIONAL MODELS

Neural systems have also been extensively modeled computationally; these include place cells (Sharp, 1991; Hartley, Burgess, Lever, Cacucci, & O'Keefe, 2000) and rat navigation (Brown & Sharp, 1995; Burgess, Donnett, Jeffery, & O'Keefe, 1997). The class of models on which I will focus will be those of reinforcement learning (RL) (Sutton & Barto, 1998), which have been used in model-based fMRI studies. RL formalizes the "law of effect" (Thorndike, 1911), which states that actions that lead to positive outcomes are more likely to be repeated. While RL models vary, they all seek to learn the value of a stimulus or action that in some way represents the reward associated with that stimulus or action. Rescorla and Wagner (1972) sought to apply this idea to classical conditioning, and devised a formula to calculate the associative strength of a conditioned stimulus after a reward. Their updated rule was interpreted as a prediction error (between the reward expected and that obtained), and was advanced by the development of a real-time, trial-by-trial temporal difference (TD) error (Sutton & Barto, 1990). The simplest TD algorithm updates the value V_s of a state, s at time, t as $V_s(t) = V_s(t-1) + \alpha[r_t + \gamma V_s(t) - V_s(t-1)]$ (Sutton & Barto, 1998), using the observed reward (r_t), a learning rate between 0 and 1 (α), and a delay discount (γ) between 0 and 1, so that delayed rewards have lower importance than immediate ones. The agent then uses these calculated values to make a decision when required, employing, for example, the softmax activation function, which converts the value of a state into a probability of action using a temperature parameter, determining the stochastic nature of the choice. This action then determines the next state the agent experiences and the reward received from the environment, and the value of the new state is then updated. These model parameters can be determined by various means, which will be discussed in the context of the application of these models to fMRI.

Neural Basis of Reinforcement Learning

It has been shown that RL algorithms can provide a good

estimation of neural activity in both animals and humans. Theories have been proposed that credit the action of the dopaminergic system and its inputs to the striatum with implementing the RL prediction error (Schultz et al., 1995); these theories are supported by single-unit recordings in monkeys (Schultz, Dayan, & Montague, 1997). Similar results were found in an fMRI experiment involving humans and a simple operant conditioning paradigm (Pagnoni, Zink, Montague, & Berns, 2002). The authors found a pattern of activity in the ventral striatum (innervated by the dopaminergic system) that showed differentiation between trials with an expected positive stimulus, and those when the stimulus was withheld. Patients with Parkinson's disease (in which striatal dopamine levels drop) have also shown difficulties when learning from feedback (Knowlton, Mangels, & Squire, 1996; Shohamy et al., 2004). These studies, however, do not fully explain what the dopamine signal represents, as it has been shown that it may represent motivation, an incentive salience, rather than an RL prediction error (Flagel et al., 2011). Despite this uncertainty, RL algorithms have been widely applied, including in the fields of spatial cognition and navigation (Foster, Morris, & Dayan, 2000; Sheynikhovich & Arleo, 2010; Gustafson & Daw, 2011).

MODEL-BASED fMRI

Model-based fMRI is a recently developed technique with the potential to uncover much more detail about how the brain carries out complex processes. All imaging-analysis methods could be considered model-based methods, in that they rely on assumptions or models of how the brain functions. Model-based fMRI, however, is a specific technique that involves using computational models to analyze how fMRI signal changes correlate with quantitative computational predictions of neural activity, rather than simply stimulus inputs and behavioral responses. This technique allows hidden variables and computational processes to be uncovered in ways not possible with traditional event-related or parametric paradigm designs. In most of the studies mentioned above, the activity reported is averaged across many trials, but by using computational models, fMRI can show not just which brain area's activity is correlated with a task but also how that brain area may carry out the task, on a trial-by-trial basis. The internal variables, such as prediction errors and state-values of RL models, calculated at each time step can be used to test different hypotheses about the possible ways the brain implements learning from reward and punishment.

Choice of Parameters

One of the main problems in the development of the RL model is that of choosing appropriate model parameters. Each of the parameters, such as the learning rate and the softmax temperature, must be chosen separately. An attractive but potentially problematic method is simply to choose parameters based on the experimental literature. However, free parameters can vary greatly among different subjects and different experimental paradigms (Kim, Shimojo, & O'Doherty, 2006; Wittmann, Daw, Seymour, &

Dolan, 2008; Daw, Gershman, Seymour, Dayan, & Dolan, 2011; Li & Daw, 2011). A popular method for estimating the free parameters is that of maximum likelihood. Optimization algorithms are available that iteratively adjust parameters to minimize the difference between the choices predicted by the model and those actually made by the subjects during the task to find the most likely parameter combinations. These algorithms are conceptually simple but are of limited use in complicated parameter spaces. Other methods (e.g., Bayesian) are conceptually more difficult and more computationally intensive, but may offer a better estimate of the parameters, and hence a better model. Once chosen, the parameters can then be used with the RL model to generate the internal variables at particular time points of the experimental task. The time series' variables are then convolved with the canonical hemodynamic function to allow for the delay between neural activity and the hemodynamic response of the neural tissue.

Hypothesis Testing-Model Comparison

As is standard in fMRI experiments, the model-predicted time series is used as a regressor against the fMRI data in a general linear model (GLM) (Friston et al., 1995). The GLM allows areas of the brain to be found where the changes in the BOLD signal have a statistically significant correlation with the model-based time series. In decision-making experiments, and fMRI experiments in general, simply finding correlated activity in a brain area doesn't show how the associated computations of the chosen model are carried out in that area. Another approach is that of model comparison, to test hypotheses of how the brain areas carry out the necessary computations for the task. Different candidate models or hypotheses may be compared to determine which model best explains the data. Often the models compared will be simple variations; in the case of RL, this could be between an on-policy TD algorithm such as SARSA (Rummery & Niranjan, 1994) or an off-policy algorithm such as Q-Learning (Watkins, 1989). Another possibility is to compare how computations are implemented more fundamentally, such as comparing model-based and model-free TD learning (Daw, Niv, & Dayan, 2005; Simon & Daw, 2011).

Simply comparing how well the different models fit the behavioral data at the maximum likelihood parameter estimates could constitute model comparison. However, this does not provide a useful answer because generally, model fit is dependent on the number of free parameters; the more free parameters there are, the better the fit. A more complex model is not necessarily better; it may just fit better to noise in the original data, and provide a worse fit to a second data set. Because of this, there are various model-comparison techniques available, such as the likelihood ratio test (Mood, Graybill, & Boes, 1974) or cross-validation (Bishop, 2006), which involves fitting the models to a subset of the data, then testing the models on the full data set. However, this method is rarely used in RL because it is difficult to split time-series data into two independent subsets (Daw, 2011). Another way of approaching model comparison is to use Bayesian methods, such as calculating

the ratio of the model evidences, known as the Bayes factor (Kass & Rafferty, 1995).

There are many methods available to compare models, but one can never be sure that the chosen model is the best available, or that a superior model will not be formulated at a later date. For this reason, a hypothesis test must be carried out to calculate the evidence in support of the null hypothesis, and whether or not it can be rejected in favor of the alternative hypothesis (the particular model to be tested).

MODEL-BASED fMRI STUDIES OF LEARNING AND DECISION MAKING

Previous fMRI studies (e.g., Pagnoni et al., 2002) found BOLD responses consistent with the prediction error (PE) in temporal-difference RL when an outcome was unexpected, but did not seek to discover whether neural activity corresponded with the predictions made by the TD algorithm throughout different stages of learning. O'Doherty, Dayan, Friston, Critchley, and Dolan (2003) used fMRI while participants took part in a Pavlovian conditioning task, and sought the neural correlates of the TD prediction error at different time points of the conditioning before, during, and after learning. The authors found activity in the ventral striatum and orbitofrontal cortex (OFC), which correlated significantly with the model-derived PE signal.

In addition to the prediction error between the reward expected and that received, it has been hypothesized that the brain might keep track of estimated rewards if previous decisions had been made differently. This would allow a distinct fictive error signal, which would further aid the organism in making future decisions. A neural correlate of this signal in the ventral caudate was found that served to modulate the behavior of subjects while they took part in an investment game (Lohrenz, McCabe, Camerer, & Montague, 2007).

A potential problem with RL is that when someone is simply learning values associated with states or actions, the higher-order structure of many tasks or environments cannot be used to make decisions. This was investigated by comparing a simple RL algorithm with a more complex computational model that incorporates the higher structure of a task carried out by participants: probabilistic reversal learning (Hampton, Bossaerts, & O'Doherty, 2006). Activity in the ventromedial prefrontal cortex (vmPFC, a region previously associated with decision making) correlated with the probability of the correct action being chosen, derived from the more complex model incorporating the structure of the task. This result is consistent with fMRI studies showing that model-generated expected-value signals associated with a stimulus are correlated with the BOLD response in various frontal cortical regions, including the vmPFC (e.g., Kim et al., 2006).

There is some uncertainty about whether the dopamine signal in the brain represents a prediction error, but there

appears to be a good concordance between RL models and neural activity in experimental animals and human subjects. These models have more recently been used to analyze fMRI data in a more detailed manner, to investigate where in the brain particular elements of a calculation are represented. This technique has provided insights into the neural basis of learning and decision making, such as the finding that activity in the ventral striatum and OFC correlate with the RL prediction error. However, until recently this has been restricted to decision making in non spatial tasks.

APPLICATION OF MODEL-BASED FMRI TO SPATIAL NAVIGATION

Since the time of the early cognitive map work, a distinction has been made between different types of spatial behavior. Blodgett and McCutchan (1947) discussed the difference between "place" and "response" learning. The former could be explained by the spatial memory encoded by place cells, and the latter could represent a simpler form of navigation, one that relied on making decisions at certain points without necessarily keeping the goal location in mind. This could well be explained by the theory of RL, and is the basis for applying model-based fMRI (using RL algorithms) to spatial navigation. Model-based fMRI has previously been applied to studies of learning and decision making; however, one study (Simon & Daw, 2011) was the first to combine model-based fMRI and VR to understand spatial navigation. Different models of how the subjects could navigate around the environment were compared—particularly whether the subjects' behavior could best be explained by a model-based RL algorithm in which the subjects used their knowledge of the structure of the environment, or by a simpler TD algorithm. In their experiment, subjects navigated in a simple 4x4 grid, with the aim of navigating toward goal locations corresponding to a monetary reward. Various models were tested, and it was found that the model-based RL algorithms fit the behavioral data much better than TD, indicating that subjects plan ahead, using a spatial map of the environment. A concurrent fMRI scan found BOLD signals within the striatum that correlated with both choice and value-related variables from the model-based RL algorithm. This is in contrast to the traditional view of the striatum being responsible for habit learning and route following. Other model-based parameters (correlated with value) were found to be correlated with activity in the medial temporal lobe and frontal cortex, concordant with previous theories about the neural basis of forward planning and internal representations of space.

This study is interesting as the first application of model-based fMRI to spatial navigation, and it begins to answer some important questions in the field. The results hint at whether people use an internal map of their environment, planning potential routes, or whether they simply follow the same paths to goals. This method could also be used to try to distinguish how the brain encodes the distance to the goal. Although it is known that the hippocampus is necessary for encoding spatial relationships, it is not known for certain whether its activity represents distance to the goal

in a Euclidian way (as the crow flies) or whether it represents the path distance, taking into account shortcuts or the distance around obstacles. Recent evidence, however, has shown that anterior hippocampal activity correlates with the Euclidian distance to the goal, and posterior hippocampal activity is linked to the path distance; which is active depends on the stage of navigation (en route, at decision points, etc.) (Howard et al., 2011). That the brain represents both Euclidian and path distance is unsurprising, as both are likely to be needed for accurate navigation in large-scale, complex environments. The value of the goal and the cost of travel are also likely to be as important as, if not more important than, the distance to be traveled when calculating paths and making decisions; model-based fMRI may provide explanations of how these variables interact in the brain.

The Simon and Daw study (2011) is a promising starting point for the method, although future studies need to be carried out to answer many of the unresolved questions within spatial navigation. The study used a highly artificial environment that, although it allows for a simple analysis, does not allow natural navigation to be investigated. To further elucidate the neural basis of decision making during navigation, a more complex, more natural environment could be used, whether with a VR environment (e.g., Hartley et al., 2003) or video recordings of natural scenes (e.g., Howard et al., 2011). Although the environment was designed to encourage a model-based strategy by incorporating dynamic rearrangement of the doors between the rooms, the subjects were always able to see the goal location above the other rooms. This makes it possible that in some of the trials, the subjects were simply trying to move closer to the goal and not thinking about the structure of the environment. If, in another experiment, the subjects were taught the environment and the goal locations prior to scanning, but then could not see them directly, they might be more likely to plan routes ahead and navigate in a more realistic manner.

LIMITATIONS OF MODEL-BASED fMRI

Although model-based fMRI is potentially a powerful technique, and has great promise in the field of spatial navigation, it is not without its limitations. Model-based fMRI is intrinsically limited by the imaging technique itself. Unlike single-unit recordings in animals, fMRI is an indirect measure of neural activity, and has low spatial and temporal resolution. As such, it can provide only an estimate of the average firing of neurons in a brain region, not the patterns of activity of individual neurons. To determine more precisely how (if at all) these algorithms are implemented in the brain, other techniques such as single-unit recordings or more-direct measures of neural activity in humans may be required, such as EEG or MEG, which could be used to uncover more accurately the time course of the activity. Another fundamental limitation of fMRI is that only a correlative, not a causal, link can be established between the neural activity and the subjects' behavior. To determine whether the region is necessary for a particular task, it

must be disrupted, either by a preexisting lesion or by the use of transcranial magnetic stimulation (Barker, Jalinous, & Freeston, 1985). As both navigation and decision making are complex processes, it is unlikely that the processes necessary for these tasks are carried out in individual brain areas. It is more likely that the computation is carried out as a dynamic pattern of activity and flow of information through many different brain areas. This is difficult to detect using simple model-based fMRI, although work has been undertaken to uncover interactions between different brain areas using techniques such as dynamic causal modeling (Friston, Harrison, & Penny, 2003). Model-based fMRI also has its own disadvantages compared to traditional fMRI, as it involves finding brain areas where the activity correlates with variables predicted by a particular model. This approach can prevent the discovery of results not expected a priori, and for this reason it is probably wise to carry out a conventional trial-based analysis of the fMRI data in conjunction with the model-based approach.

Rather than just using model-based fMRI or comparing results with other techniques, a similar model-based analysis could be applied to any physiological measure that correlates with behavior. RL models could be adapted to carry out model-based analyses of imaging data from other, complementary techniques such as EEG/MEG or measures such as eye tracking. These methods could be formally combined—for example, simultaneous EEG-fMRI recording (see Laufs, Daunizeau, Carmichael, & Kleinschmidt [2008] for a review)—to potentially provide an insight into the computational processes carried out by the brain during navigation at a high spatial and temporal resolution.

The most important limitation of any model-based analysis is the assumptions it requires. Model-based fMRI requires many steps, from constructing the models and designing the experiment to collecting and analyzing the data. In all this it is easy to forget that the whole technique relies on an assumption that the brain reduces a very complicated problem to a few simple steps with particular variables. It is important to remember that this may be a flawed construct, and just because the analysis gives an appealing answer, that does not mean it is necessarily true. Any evidence this method provides must be interpreted in the light of the rest of the experimental literature, and supported by results obtained by other methods.

The application of model-based fMRI to spatial navigation research is promising, although only one study has yet been performed (which was designed to study decision making in a navigation paradigm, rather than navigation itself). This method has the potential, however, to reveal how humans use internal models of their environment, how they assign value to parts of their environment, and how they use this information to make decisions and navigate accurately. But to accomplish this, care must be taken to design tasks that will allow these variables to be investigated, while ensuring that the navigation paradigm corresponds well to real-world tasks. Model-based fMRI has many limitations, most

of which are simply the limitations of the imaging modality itself; they can be overcome, at least in part, by combining or comparing results from model-based fMRI with results from other methods. However, model-based fMRI rests entirely on the validity of the models chosen for the analysis, and this must be kept in mind when one is interpreting results.

CONCLUSION

Much has been learned about spatial memory, from single-unit recordings in animals to sophisticated imaging studies in humans. RL algorithms have helped our understanding of how we and other animals learn about our environment and make decisions. The application of these models in model-based fMRI results in a particularly powerful technique, allowing researchers to detect where in the brain specific elements of particular computations are carried out. The study by Simon and Daw (2011) is a promising starting point for the application of the method to spatial navigation; however, as the authors acknowledge, their study was more like others designed to investigate decision making rather than spatial navigation. To investigate the unanswered questions in navigation, such as how we make decisions and use models of our environment, the method needs to be improved, advanced, and perhaps supplemented with model-based analyses of other techniques, such as EEG/MEG.

Because model-based fMRI relies on the validity of applying RL models to spatial navigation, this must also be investigated. Experiments in animals may provide a method for doing just this. Selective inactivation, both in time and space, of areas thought to be involved in RL-related processes as an animal learns and navigates within an environment could be used to investigate the validity of applying models originally developed to explain learning in conditioning paradigms to spatial navigation by humans in complex environments.

Corresponding Author

Adam Tyson (adam.tyson.09@ucl.ac.uk), Faculty of Mathematical and Physical Sciences, University College London, Gower Street, London, WC1E 6BT, United Kingdom.

Conflict of Interest Disclosure

The author has completed and submitted the ICMJE Form for Disclosure of Potential Conflicts of Interest. No conflicts were noted.

Acknowledgments

The author thanks Dr. Hugo Spiers for his helpful suggestions in the writing of this article.

References

- Aguirre, G. K., Detre, J. A., Alsup, D. C., & D'Esposito, M. (1996). The parahippocampus subserves topographical learning in man. *Cerebral Cortex*, 6(6), 823–829.
- Aimone, J. B., Wiles, J., & Gage, F. H. (2006). Potential role for adult neurogenesis in the encoding of time in new memories. *Nature Neuroscience*, 9(6), 723–727.
- Barker, A. T., Jalinous, R., & Freeston, I. L. (1985). Non-invasive magnetic stimulation of human motor cortex. *Lancet*, 1(8437), 1106–1107.
- Becker, S. (2005). A computational principle for hippocampal learning and neurogenesis. *Hippocampus*, 15(6), 722–738.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. New York, NY: Springer.
- Blodgett, H. C. (1929). The effect of the introduction of reward upon the maze performance of rats. *University of California Publications in Psychology*, 4, 113–134.
- Blodgett, H. C., & McCutchan, K. (1947). Place versus response learning in the simple T-maze. *Journal of Experimental Psychology*, 37(5), 412–422.
- Brown, M. A., & Sharp, P. E. (1995). Simulation of spatial learning in the Morris water maze by a neural network model of the hippocampal formation and nucleus accumbens. *Hippocampus*, 5(3), 171–188.
- Burgess, N., Donnett, J. G., Jeffery, K. J., & O'Keefe, J. (1997). Robotic and neuronal simulation of the hippocampus and rat navigation. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences*, 352(1360), 1535–1543.
- Burgess, N., Jackson, A., Hartley, T., & O'Keefe, J. (2000). Predictions derived from modelling the hippocampal role in navigation. *Biological Cybernetics*, 83(3), 301–312.
- Corlett, P. R., Aitken, M. R., Dickinson, A., Shanks, D. R., Honey, G. D., Honey, R. A., . . . Fletcher, P. C. (2004). Prediction error during retrospective reevaluation of causal associations in humans: fMRI evidence in favor of an associative model of learning. *Neuron*, 44(5), 877–888.
- Cornwell, B. R., Johnson, L. L., Holroyd, T., Carver, F. W., & Grillon, C. (2008). Human hippocampal and parahippocampal theta during goal-directed spatial navigation predicts performance on a virtual Morris water maze. *Journal of Neuroscience*, 28(23), 5983–5990.
- Daw, N. D. (2011). Trial-by-trial data analysis using computational models. In M. R. Delgado, E. A. Phelps, & T. W. Robbins (Eds.), *Decision making, affect, and learning: Attention and performance XXIII* (pp. 3–38). Oxford, England: Oxford University Press.
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6), 1204–1215.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 8(12), 1704–1711.
- de Araujo, D. B., Baffa, O., & Wakai, R. T. (2002). Theta oscillations and human navigation: A magnetoencephalography study. *Journal of Cognitive Neuroscience*, 14(1), 70–78.
- Doeller, C. F., Barry, C., & Burgess, N. (2010). Evidence for grid cells in a human memory network. *Nature*, 463(7281), 657–U687.
- Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*, 425(6954), 184–188.
- Flagel, S. B., Clark, J. J., Robinson, T. E., Mayo, L., Czuj, A., Willuhn, I., . . . Akil, H. (2011). A selective role for dopamine in stimulus-reward learning. *Nature*, 469(7328), 53–U63.
- Foster, D. J., Morris, R. G. M., & Dayan, P. (2000). A model of hippocampally dependent navigation, using the temporal difference learning rule. *Hippocampus*, 10(1), 1–16.
- Friston, K. J., Harrison, L., & Penny, W. (2003). Dynamic causal modelling. *Neuroimage*, 19(4), 1273–1302.
- Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J. P., Frith, C. D., & Frackowiak, R. S. J. (1995). Statistical parametric maps in functional imaging: A general linear approach. *Human Brain Mapping*, 2, 189–210.
- Gustafson, N. J., & Daw, N. D. (2011). Grid cells, place cells, and geodesic generalization for spatial reinforcement learning. *PLoS Computational Biology*, 7(10): e1002235.
- Hafting, T., Fyhn, M., Molden, S., Moser, M. B., & Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052), 801–806.
- Hampton, A. N., Bossaerts, P., & O'Doherty, J. P. (2006). The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. *Journal of Neuroscience*, 26(32), 8360–8367.
- 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(4), 369–379.
- Hartley, T., Maguire, E. A., Spiers, H. J., & Burgess, N. (2003). The well-worn route and the path less traveled: Distinct neural bases of route following and wayfinding in humans. *Neuron*, 37(5), 877–888.
- Henderson, V. W., Mack, W., & Williams, B. W. (1989). Spatial disorientation in Alzheimer's disease. *Archives of Neurology*, 46(4), 391–394.
- Howard, L. R., Yu, Y., Mill, R. D., Morrison, L. C., Knight, R., Loftus, M., . . . Spiers, H. J. (2011, November). *Human hippocampus encodes Euclidean distance and future path to goals during real-world navigation*. Paper presented at the 41st Annual Meeting of the Society for Neuroscience, Washington, DC.
- Kahana, M. J., Sekuler, R., Caplan, J. B., Kirschen, M., & Madsen, J. R. (1999). Human theta oscillations exhibit task dependence during virtual maze navigation. *Nature*, 399(6738), 781–784.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–795.
- Kee, N., Teixeira, C. M., Wang, A. H., & Frankland, P. W. (2007). Preferential incorporation of adult-generated granule cells into spatial memory networks in the dentate gyrus. *Nature Neuroscience*, 10(3), 355–362.
- Kim, H., Shimojo, S., & O'Doherty, J. P. (2006). Is avoiding an aversive outcome

- rewarding? Neural substrates of avoidance learning in the human brain. *PLoS Biology*, 4(8), 1453–1461.
- Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996). A neostriatal habit learning system in humans. *Science*, 273(5280), 1399–1402.
- Laufs, H., Daunizeau, J., Carmichael, D. W., & Kleinschmidt, A. (2008). Recent advances in recording electrophysiological data simultaneously with magnetic resonance imaging. *Neuroimage*, 40(2), 515–528.
- Li, J., & Daw, N. D. (2011). Signals in human striatum are appropriate for policy update rather than value prediction. *Journal of Neuroscience*, 31(14), 5504–5511.
- Lohrenz, T., McCabe, K., Camerer, C. F., & Montague, P. R. (2007). Neural signature of fictive learning signals in a sequential investment task. *Proceedings of the National Academy of Sciences of the United States of America*, 104(22), 9493–9498.
- Maguire, E. A., Burgess, N., Donnett, J. G., Frackowiak, R. S., Frith, C. D., & O'Keefe, J. (1998). Knowing where and getting there: A human navigation network. *Science*, 280(5365), 921–924.
- Maguire, E. A., Burke, T., Phillips, J., & Staunton, H. (1996). Topographical disorientation following unilateral temporal lobe lesions in humans. *Neuropsychologia*, 34(10), 993–1001.
- Maguire, E. A., Gadian, D. G., Johnsrude, I. S., Good, C. D., Ashburner, J., Frackowiak, R. S., & Frith, C. D. (2000). Navigation-related structural change in the hippocampi of taxi drivers. *Proceedings of the National Academy of Sciences of the United States of America*, 97(8), 4398–4403.
- Maguire, E. A., Woollett, K., & Spiers, H. J. (2006). London taxi drivers and bus drivers: A structural MRI and neuropsychological analysis. *Hippocampus*, 16(12), 1091–1101.
- Mood, A. M., Graybill, F. A., & Boes, D. C. (1974). *Introduction to the theory of statistics* (3rd ed.). New York, NY: McGraw-Hill.
- Morris, R. G., Garrod, P., Rawlins, J. N., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*, 297(5868), 681–683.
- O'Doherty, J. P., Dayan, P., Friston, K., Critchley, H., & Dolan, R. J. (2003). Temporal difference models and reward-related learning in the human brain. *Neuron*, 38(2), 329–337.
- O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34(1), 171–175.
- O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford, England: Oxford University Press.
- Pagnoni, G., Zink, C. F., Montague, P. R., & Berns, G. S. (2002). Activity in human ventral striatum locked to errors of reward prediction. *Nature Neuroscience*, 5(2), 97–98.
- Rauchs, G., Orban, P., Baiteau, E., Schmidt, C., Degueldre, C., Luxen, A., . . . Peigneux, P. (2008). Partially segregated neural networks for spatial and contextual memory in virtual navigation. *Hippocampus*, 18(5), 503–518.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York, NY: Appleton-Century-Crofts.
- Rummery, G. A., & Niranjan, M. (1994). *On-line Q-learning using connectionist systems*. Technical Report CUED/F-INFENG/TR 166, Engineering Department. Cambridge, England: Cambridge University.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*, 275(5306), 1593–1599.
- Schultz, W., Romo, R., Ljungberg, T., Mirenowicz, J., Hollerman, J. R., & Dickinson, A. (1995). Reward-related signals carried by dopamine neurons. In J. C. Houk, J. L. Davis, & D. G. Beiser (Eds.), *Models of information processing in the basal ganglia* (pp. 233–248). Cambridge, MA: MIT Press.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery and Psychiatry*, 20(1), 11–21.
- Sharp, P. E. (1991). Computer simulation of hippocampal place cells. *Psychobiology*, 19, 103–115.
- Sheynikhovich, D., & Arleo, A. (2010). A reinforcement learning approach to model interactions between landmarks and geometric cues during spatial learning. *Brain Research*, 1365, 35–47.
- Shohamy, D., Myers, C. E., Grossman, S., Sage, J., Gluck, M. A., & Poldrack, R. A. (2004). Cortico-striatal contributions to feedback-based learning: Converging data from neuroimaging and neuropsychology. *Brain*, 127, 851–859.
- Simon, D. A., & Daw, N. D. (2011). Neural correlates of forward planning in a spatial decision task in humans. *Journal of Neuroscience*, 31(14), 5526–5539.
- Spiers, H. J., Burgess, N., Hartley, T., Vargha-Khadem, F., & O'Keefe, J. (2001). Bilateral hippocampal pathology impairs topographical and episodic memory but not visual pattern matching. *Hippocampus*, 11(6), 715–725.
- Spiers, H. J., & Maguire, E. A. (2006). Thoughts, behaviour, and brain dynamics during navigation in the real world. *Neuroimage*, 31(4), 1826–1840.
- Spiers, H. J., & Maguire, E. A. (2007). A navigational guidance system in the human brain. *Hippocampus*, 17(8), 618–626.
- Stone, S. S., Teixeira, C. M., DeVito, L. M., Zaslavsky, K., Josselyn, S. A., Lozano, A. M., & Frankland, P. W. (2011). Stimulation of entorhinal cortex promotes adult neurogenesis and facilitates spatial memory. *Journal of Neuroscience*, 31(38), 13469–13484.
- Suthana, N., Haneef, Z., Stern, J., Mukamel, R., Behnke, E., Knowlton, B., & Fried, I. (2012). Memory enhancement and deep-brain stimulation of the entorhinal area. *New England Journal of Medicine*, 366(6), 502–510.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- Sutton, R. S., & Barto, A. G. (1990). Time-derivative models of Pavlovian reinforcement. In M. Gabriel & J. Moore (Eds.), *Learning and computational neuroscience: Foundations of adaptive networks* (pp. 497–537). Cambridge, MA: MIT Press.
- Taube, J. S., Muller, R. U., & Ranck, J. B., Jr. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats: 1. Description and quantitative analysis. *Journal of Neuroscience*, 10(2), 420–435.
- Thorndike, E. L. (1911). *Animal intelligence: An experimental study of the associative processes in animals*. New York, NY: Macmillan.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55(4), 189–208.
- Tolman, E. C., Ritchie, B. F., & Kalish, D. (1946). Studies in spatial learning: 1. Orientation and the short-cut. *Journal of Experimental Psychology*, 36, 13–24.
- Watkins, C. (1989). *Learning from delayed rewards* (Unpublished doctoral dissertation). University of Cambridge, Cambridge, England.
- Wittmann, B. C., Daw, N. D., Seymour, B., & Dolan, R. J. (2008). Striatal activity underlies novelty-based choice in humans. *Neuron*, 58(6), 967–973.
- Xu, J., Evensmoen, H. R., Lehn, H., Pintzka, C. W., & Häberg, A. K. (2010). Persistent posterior and transient anterior medial temporal lobe activity during navigation. *Neuroimage*, 52(4), 1654–1666.