

NEURAL REPRESENTATIONS OF A REAL-WORLD ENVIRONMENT

Lindsay K. Vass

A DISSERTATION

in

Neuroscience

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2014

Supervisor of Dissertation

Russell A. Epstein

Professor, Psychology

Graduate Group Chairperson

Joshua I. Gold, Professor, Neuroscience

Dissertation Committee

Yale E. Cohen, Associate Professor, Otorhinolaryngology

Joseph W. Kable, Assistant Professor, Psychology

Isabel A. Muzzio, Assistant Professor, Psychology

Kenneth A. Norman, Professor, Psychology

NEURAL REPRESENTATIONS OF A REAL-WORLD ENVIRONMENT

COPYRIGHT

2014

Lindsay Katherine Vass

ACKNOWLEDGMENTS

I am extraordinarily grateful to a number of people whose support has helped make this dissertation possible.

First, I must thank my advisor, Russell Epstein. Over the past six years, he has shown by example what it means to be a great scientist. He has helped me to become a better speaker, a better writer, and a better thinker. Most importantly, he has always been a font of positive support, and his feedback, complete with dozens of exclamation points, has kept me going when experiments went awry or when the scientific truth seemed just out of my grasp. He also introduced me to the Old Fashioned, and for that, I am eternally grateful. I also thank my committee—Yale Cohen, Joe Kable, and Isabel Muzzio—whose diverse expertise helped me to ponder spatial cognition from a variety of perspectives.

Second, I have had the pleasure of doing science alongside a lab full of exceptional people who not only work hard, but play hard too. This ethos is best captured by our adventures at Vision Sciences Society, and although it will probably never be my favorite conference, the memories will always have a privileged place in my hippocampus. I also have to specifically thank Mary Smith, who welcomed me to the lab and taught me VoxBo (although it's no longer officially supported, I won't let it die just yet!) and Steven Marchette, who has spent many hours commiserating with me over the challenges of studying spatial cognition and many more reminding me what a privilege it is to do what we do.

Third, I have been supported by the fantastic (and now award-winning!) Neuroscience Graduate Group. It is invigorating to be a part of such a tight-knit and supportive community, and that atmosphere is due in no small part to the three people who have chaired NGG over my graduate school career: Mikey Nusbaum, Rita Balice-Gordon, and Josh Gold. And of course, without NGG, I wouldn't have met such an incredible group of friends. "Girls Night," organized by one of my dearest friends,

Vanessa Troiani, kept me (mostly) sane throughout graduate school and is something I will certainly miss.

Fourth, I have to thank my family, who have encouraged and supported my academic endeavors my entire life. My parents, Mike and Karen Morgan, scrimped and saved to send me to private school (sorry for not always appreciating this) and CTY, also known as “nerd camp,” where I finally met likeminded people whose ideal summer vacation consisted of seven hours of math a day. I thank them, along with my sisters, Mandy and Audrey, for all the love they have given me over the years.

Finally, I thank my husband, Brian Vass, who has been my very best friend for these last nine years. Words cannot express how thankful I am for his never-ending love and encouragement.

ABSTRACT

NEURAL REPRESENTATIONS OF A REAL-WORLD ENVIRONMENT

Lindsay K. Vass

Russell A. Epstein

The ability to represent the spatial structure of the environment is critical for successful navigation. Extensive research using animal models has revealed the existence of specialized neurons that appear to code for spatial information in their firing patterns. However, little is known about which regions of the human brain support representations of large-scale space. To address this gap in the literature, we performed three functional magnetic resonance imaging (fMRI) experiments aimed at characterizing the representations of locations, headings, landmarks, and distances in a large environment for which our subjects had extensive real-world navigation experience: their college campus. We scanned University of Pennsylvania students while they made decisions about places on campus and then tested for spatial representations using multivoxel pattern analysis and fMRI adaptation. In Chapter 2, we tested for representations of the navigator's current location and heading, information necessary for self-localization. In Chapter 3, we tested whether these location and heading representations were consistent across perception and spatial imagery. Finally, in Chapter 4, we tested for representations of landmark identity and the distances between landmarks. Across the three experiments, we observed that specific regions of medial temporal and medial parietal cortex supported long-term memory representations of navigationally-relevant spatial information. These results serve to elucidate the functions of these regions and offer a framework for understanding the relationship between spatial representations in the medial temporal lobe and in high-level visual regions. We discuss our findings in the context of the broader spatial cognition literature, including implications for studies of both humans and animal models.

TABLE OF CONTENTS

ACKNOWLEDGMENTS	III
ABSTRACT	V
LIST OF ILLUSTRATIONS.....	VIII
CHAPTER 1: INTRODUCTION	1
1.1 Spatial Representations.....	1
1.2 Evidence from neurophysiological recordings.....	3
1.3 Evidence from neuropsychology.....	6
1.4 Evidence from neuroimaging	8
CHAPTER 2: ABSTRACT REPRESENTATIONS OF LOCATION AND FACING DIRECTION IN THE HUMAN BRAIN	14
2.1 Abstract	14
2.2 Introduction.....	14
2.3 Materials and Methods	16
2.4 Results	24
2.5 Discussion	33
CHAPTER 3: NEURAL REPRESENTATIONS UNDERLYING REAL-WORLD SPATIAL MEMORY RETRIEVAL	48
3.1 Abstract	48
3.2 Introduction.....	49
3.3 Materials and Methods	51
3.4 Results	62
3.5 Discussion	70
CHAPTER 4: DISTANCES BETWEEN REAL-WORLD LOCATIONS ARE REPRESENTED IN THE HUMAN HIPPOCAMPUS	85

4.1 Abstract	85
4.2 Introduction.....	85
4.3 Materials and Methods	87
4.4 Results	92
4.5 Discussion	99
CHAPTER 5: GENERAL DISCUSSION AND CONCLUSIONS	112
5.1 Relevance to neuropsychological literature	113
5.2 Relevance to human behavioral literature	114
5.3 Relevance to computational modeling	115
5.4 Relevance to neurophysiological studies.....	117
5.5 Relevance to neuroimaging studies	119
5.6 Representations in PPA, RSC, and Hippocampus.....	120
5.7 Conclusions	124
BIBLIOGRAPHY	126

LIST OF ILLUSTRATIONS

- 2.1 Meta-analysis of brain regions activated during studies of navigation.
- 2.2 Map and example stimuli used in the experiment.
- 2.3 Coding of spatial quantities in multivoxel patterns in PPA and RSC.
- 2.4 Coding of spatial quantities in multivoxel patterns in the medial temporal lobe.
- 2.5 Whole-brain searchlight analyses of multivoxel coding of view, location, and direction information.
- 2.6 Direction tuning curves for right presubiculum.
- 2.7 Photographs of 32 views and 17 exemplar photographs for one view.
- 3.1 Experimental design and task.
- 3.2 Behavioral priming.
- 3.3 Multivoxel pattern correlations in PPA and RSC.
- 3.4 Multivoxel pattern correlations in MTL ROIs.
- 3.5 Searchlight analysis for heading.
- 3.6 Location representation in RSC, PPA, and ERC.
- 3.7 Heading representation in RSC, PPA, and ERC.
- 4.1 Examples of stimuli and map showing the locations of the 10 landmarks on the University of Pennsylvania campus.
- 4.2 Whole-brain analysis for landmark adaptation.
- 4.3 Distance-related adaptation in the human brain.
- 4.4 Decoding of landmark identity using MVPA.
- 4.5 Whole-brain (searchlight) analysis.
- 4.6 Examples of the twenty-two images of one landmark, Huntsman Hall.
- 4.7 Decoding of individual landmarks using MVPA.
- 5.1 Model of medial temporal and medial parietal representations.

CHAPTER 1: INTRODUCTION

1.1 Spatial Representations

All mobile organisms need to move through the world in search of food, a mate, and safe haven from predators. As a consequence, organisms that can faithfully represent the environmental space around them and store this information in memory will be conferred an evolutionary advantage. Animals as humble as the honeybee have perfected this ability and for many years great minds have wondered how. Specifically, how does one transform the continuous perceptual experience of the external world into a coherent internal representation that captures the spatial structure of the environment? The basic neuroscience approach to this question has been to search for neural signals that correlate with spatial properties. Thus, an important first question is what features of space would be useful to represent? At a very basic level, two pieces of information are critical. First, one would want to represent the identities of discrete places, the sensory features that reliably distinguish one place from another. This type of representation would afford an animal the ability to make more optimal decisions like seeking food from the bushes by the river rather than the bushes by the lion's den. Second, one would want to represent the spatial relationships between places, such as the distances and angles between them, or the routes that link them. This knowledge allows for efficient goal-driven behavior rather than relying on opportunistic wandering. Because all locations are relative, these spatial relationships must be represented within a particular *reference frame*, which is characterized by a *reference point*, or origin, and a *reference direction*, or axis. There are two general classes of reference frames that are used for navigation. The first is an *egocentric* reference frame, in which distances and directions are coded with respect to the observer (e.g., "the bench is 3 meters to my right"). In an egocentric reference frame, the reference point is typically the observer and the reference direction is the direction the navigator is facing, i.e., the navigator's *heading*. The second is an *allocentric* reference frame, in which distances

and directions are coded with respect to external referents (e.g., “the bench is 3 meters north of the trash can”). In this case, the reference point may be any arbitrary location, such as a salient visible landmark. Likewise, the reference direction could be any arbitrary direction, and is often specified with respect to a distal landmark (“towards the mountains”) or the earth’s magnetic field (“north”). Because the reference direction is not tied to the navigator, in order to use an allocentric reference frame, a navigator must also have a way to represent their current heading, the angle between the navigator’s axis of orientation and the reference direction.

One of the central questions in spatial cognition is which reference frames are used to represent space. In 1948, Tolman famously proposed that rodents navigated via an allocentric “cognitive map” because they were able to take a novel shortcut through a maze to find a food reward. To this day, there is continued debate about whether humans use allocentric representations at all, with some suggesting that behavior can be explained with a dynamically updated egocentric representation that re-calculates the self-object vectors each time the navigator moves (Wang and Spelke, 2002; Wang, 2012). However, most models assume that a navigator can use both egocentric and allocentric reference frames and can flexibly transform information in one reference frame to the other (Gallistel, 1990; Klatzky, 1998; Sholl, 2001; Mou et al., 2004; Byrne et al., 2007; Whitlock et al., 2008). What differentiates many of these multiple-representation models are the features of the environment that the allocentric reference frame is anchored to and the circumstances under which representations are transformed.

In the following sections, I focus on allocentric representations of space and review the neural evidence for their existence. I first discuss evidence from recordings of individual neurons, which are usually performed in rodents, but are occasionally performed in monkeys and humans. I then describe evidence from human neuropsychological patients, who have had focal brain lesions that disrupted their

ability to navigate. Finally, I discuss evidence of neural representations of space in the intact human brain, as assessed by neuroimaging.

1.2 Evidence from neurophysiological recordings

Neurophysiological recordings have provided the clearest evidence for a neural representation of space. In these studies, typically performed in rodents, an electrode is implanted into the animal's brain, which allows the researcher to record the neural activity of a small number of neurons while the animal freely explores an environment. The seminal finding that launched the field was the observation that individual neurons in the hippocampus fired action potentials when the animal occupied a particular location in the environment (O'Keefe and Dostrovsky, 1971; O'Keefe, 1976). These "place cells," which have since been identified in humans (Ekstrom et al., 2003; Miller et al., 2013), were hypothesized to be the neural basis of the cognitive map (O'Keefe and Nadel, 1978). Many studies since then have sought to identify the precise cues that contribute to place cell firing. These studies, which typically record place cells before and after manipulation of an environmental feature, have shown that many different cues can drive place cell responses, including proximal and distal landmarks (O'Keefe and Conway, 1978; Shapiro et al., 1997; Knierim, 2002), environmental boundaries (O'Keefe and Burgess, 1996), and self-motion (Quirk et al., 1990; for a review, see Knierim and Hamilton, 2011). In sum, place cells appear to code for the position of the navigator in a way that is referenced to external features of the environment.

In the years since that initial discovery, many other types of spatial cells have been observed in the hippocampus and surrounding medial temporal lobe (MTL). Here I focus on two of the most well-characterized cell types, head direction (HD) cells and grid cells, which are hypothesized to represent two key features of an allocentric representation: the navigator's heading and the map's coordinate system. HD cells were first identified in the dorsal presubiculum, a structure near the hippocampus, and fire whenever the animal's head is oriented in a particular allocentric direction (e.g., north) independent of the animal's position in space (Taube et al., 1990b). Each cell

has its own preferred direction, and simultaneous recordings of multiple HD cells have shown that the relative firing preferences across neurons is constant across environments (e.g., if neuron A prefers 0° and neuron B prefers 90° in a given environment, their preferences will continue to be offset by 90° in other environments even if the absolute direction preferences change), suggesting that they are organized in a highly interconnected circuit (Taube et al., 1990a). Thus, like a compass, a population of HD cells represents the animal's current perceived heading in the environment. However, unlike a compass, these cells do not respond to differences in the earth's magnetic field; rather, they are driven by self-motion cues originating in the vestibular system and are calibrated by distal visual landmarks in the environment (Winter and Taube, 2014). After their initial discovery in the presubiculum, HD cells have since been found in a number of interconnected regions including parasubiculum (Boccarda et al., 2010), entorhinal cortex (Sargolini et al., 2006), retrosplenial cortex (Chen et al., 1994; Cho and Sharp, 2001) and thalamus (Taube, 1995), with cells in each region exhibiting largely similar firing characteristics (Winter and Taube, 2014). These cells have also been identified in monkeys (Robertson et al., 1999) and putative HD cells have been observed in humans (Jacobs et al., 2010). Although work is still in progress to determine the precise contributions of each region to the animal's sense of direction, one notable observation is that cells in the presubiculum seem to be particularly important for calibrating directional preferences on the basis of visual landmarks (Yoder et al., 2011). Thus, this region may play an important role in relaying perceptual information to the long-term spatial memory system.

The third cell type are grid cells, which fire in regularly-spaced locations on a hexagonal lattice that tiles the entire environment (Hafting et al., 2005). These cells are found in entorhinal cortex, the primary input structure to the hippocampus, as well as adjacent pre- and parasubiculum (Sargolini et al., 2006). Recently, they have also been observed in humans (Jacobs et al., 2013). Because of the exquisite regularity of their firing fields, these cells are believed to provide the coordinate system of the cognitive

map, a feature which would allow for the calculation of metric values, such as the distance between landmarks. For example, to measure distance traveled, a navigator could simply count the number of grid fields that he or she has passed through (see Kubie and Fenton, 2012 for a model of vector calculation using grid cells). Recent work has shown that different populations of grid cells have different amounts of spacing between the grid fields (Stensola et al., 2012). As a consequence, this affords the ability to represent environments of different sizes and at different spatial resolutions (Jeffery, 2013).

In summary, neurophysiological recordings have provided evidence for an allocentric map of space in the MTL, with place cells and HD cells coding the current location and heading of the navigator respectively, and grid cells coding the coordinate system of the representation. Additional cell types not discussed here code for features of the environment itself, such as specific regions of visible space (Georges-Francois et al., 1999; Rolls, 1999; Ekstrom et al., 2003), or the navigator's position relative to such features, such as boundaries (Solstad et al., 2008; Lever et al., 2009) or landmarks (Deshmukh and Knierim, 2013). These types of representations appear to be largely conserved across mammalian species, as cellular recordings from human MTL have revealed remarkably similar cell types to those observed in rodents (Ekstrom et al., 2003; Jacobs et al., 2010; Jacobs et al., 2013; Miller et al., 2013; but see Jacobs, 2014). Much more work is needed to fully characterize the neurons that represent space in humans. Due to the invasive nature of electrophysiological recordings, these studies can only be carried out in human patients who have had electrodes implanted for surgical purposes, typically to identify the epileptogenic focus in cases of drug-resistant epilepsy (Jacobs and Kahana, 2010). As a consequence, these data are necessarily challenging to acquire. However, we can also interrogate human spatial representations indirectly by studying patients with focal brain lesions who have lost their ability to navigate. I describe evidence from these case studies in the next section.

1.3 Evidence from neuropsychology

Prior to the advent of neuroimaging techniques, the primary evidence for the neural localization of human mental functions was obtained by observing patients with focal brain lesions. The logic is that if a patient is no longer able to execute a particular behavior after the neurological insult, then the part of the brain that was damaged is likely to be involved in generating that behavior. Patients that lose the ability to navigate are described as having *topographical disorientation* (TD), and numerous case reports over the years have revealed that different patients exhibit distinct patterns of navigational impairments, which are associated with damage to different regions of the brain (Aguirre and D'Esposito, 1999).

The first variant of TD is *landmark agnosia*, an inability to perceive and recognize prominent environmental features such as buildings and landscapes. This deficit is typically associated with damage to the parahippocampal and/or lingual gyrus, structures in the MTL. Strikingly, these patients generally have preserved spatial knowledge, as they are able to draw maps of familiar environments and describe routes between familiar places (Pallis, 1955; Epstein et al., 2001; Takahashi and Kawamura, 2002; Mendez and Chierri, 2003). Thus, the navigational deficit appears to arise from an inability to analyze the visual scene and extract the relevant features. As one patient described it, "In my mind's eye, I know exactly where places are, what they look like...It's when I'm out that the trouble starts. My reason tells me I must be in a certain place and yet I don't recognize it" (Pallis, 1955).

This is quite a different behavioral profile from the second variant of TD, *heading disorientation*. Patients with this disorder can recognize landmarks, but cannot retrieve the directional relationships between them. For example, a taxicab driver suffered a cerebral hemorrhage while driving and "suddenly lost his understanding of the route to his destination. As he could quickly recognize the buildings and landscape around him, he was able to determine his current location. However, he could not determine in which direction to proceed" (Takahashi et al., 1997). In other words, he appeared to

have lost his sense of direction. Unlike landmark agnosics, these patients are often unable to draw maps or describe routes, suggesting that they cannot access spatial information either via perception or mental imagery (Takahashi et al., 1997; Luzzi et al., 2000; Tamura et al., 2007; Osawa et al., 2008). Although there have been reports of heading disorientation after damage to right parahippocampal cortex (Alsaadi et al., 2000; Luzzi et al., 2000; Caglio et al., 2011), this behavioral profile most commonly arises after damage to medial parietal or retrosplenial cortex (Takahashi et al., 1997; Alsaadi et al., 2000; Tamura et al., 2007; Osawa et al., 2008; Hashimoto et al., 2010). Based on the majority of reports, it appears that medial parietal cortex is involved in representing the directional relationships between landmarks. It may also be involved in *self-localization*, the ability to identify one's current position and heading. Although this ability is rarely explicitly tested, one patient with damage to right medial parietal cortex exhibited a selective self-localization impairment despite being able to draw maps and describe routes (Suzuki et al., 1998). When viewing photographs of her home, she was unable to infer the position of the photographer and when actively navigating, she was unable to indicate her current position on a map. Another patient also showed a similar deficit and was unable to indicate her location on a map or on a miniature model of the environment (Katayama et al., 1999). In sum, whereas parahippocampal cortex appears to be involved in processing the visual scene, retrosplenial and medial parietal cortex appear to be involved in associating that scene with the appropriate spatial information.

Based on the neurophysiological data, one would expect that damage to the hippocampus should also result in TD. Although the hippocampus does appear to be necessary for forming new spatial memories (but see Corkin, 2002), it is not required for navigating environments learned prior to injury (Teng and Squire, 1999; Rosenbaum et al., 2000; Maguire et al., 2006). Patients KC, HM, TT, and EP, all of whom sustained a nearly complete loss of their bilateral hippocampi and were profoundly amnesic, retained spatial knowledge of their childhood neighborhood. However, there is some

evidence that their spatial representations were not as rich or detailed as neurologically intact controls. KC's sketch maps contained fewer landmarks than controls and while he could recognize photographs of the most salient neighborhood landmarks, like a school or shopping center, he failed to recognize individual houses (Rosenbaum et al., 2000). TT, the only patient tested during active navigation (in a video game replicate of London), could successfully navigate when on main roads, but became lost on smaller streets (Maguire et al., 2006). Thus, structures outside the hippocampus such as retrosplenial cortex may be able to support coarse navigation and spatial knowledge, but the hippocampus appears to be necessary for the retrieval of finely detailed spatial information. To better understand the precise roles these regions play in navigation, I now turn to evidence from neuroimaging studies, which have afforded the ability to understand the neural basis of mental functions in the intact brain.

1.4 Evidence from neuroimaging

Neuroimaging techniques including positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) provide the opportunity to examine which regions in the intact brain are engaged by navigation tasks. A meta-analysis of 13 neuroimaging studies of virtual or imagined navigation indicates a core network of regions are activated, including hippocampus, parahippocampal cortex, and retrosplenial cortex, consistent with the regions implicated by the neuropsychological literature (Spreng et al., 2009). What precise roles do each of these regions play? Navigation is an inherently complex behavior, involving perception of environmental and idiothetic cues, representation of spatial quantities like position and heading, and various computations on these quantities, such as transforming between egocentric and allocentric reference frames (Wolbers and Hegarty, 2010). As such, the brain regions identified by the navigation meta-analysis could contribute to any or all of these processes. To begin to delineate the roles of each region, I will first consider evidence from studies that have examined the level of activation in these regions during various

spatial tasks. I will then describe evidence for the possible representations supported by these regions during these tasks.

Although hippocampus, parahippocampal cortex, and retrosplenial cortex are all engaged by navigation tasks, only parahippocampal and retrosplenial cortex are typically activated by spatial memory tasks that do not require active navigation (Committeri et al., 2004; Rosenbaum et al., 2004; Epstein et al., 2007; Galati et al., 2010; Schinazi and Epstein, 2010). There is strong evidence that these two regions are involved in the perception of spatial information during navigation. Both regions are more active when subjects view images of buildings or objects previously encountered at navigationally important locations (i.e., intersections; Janzen 2004; Schinazi 2010) and are even strongly activated when subjects passively view navigationally relevant stimuli (i.e., scenes) in the absence of any spatial task (Epstein and Kanwisher, 1998). This increased response to scenes is quite selective, as these regions respond only weakly to images of objects and not at all to images of faces. When functionally defined based on their preference for scenes, these regions are referred to as the parahippocampal place area (PPA) and the retrosplenial complex (RSC). Although defined based on responses to visual stimuli, there is some evidence that these regions are not strictly visual, as both PPA and RSC are activated in blind subjects when they haptically explore scenes made of Lego blocks (Wolbers et al., 2011). This finding suggests that the codes in these regions are at least partially spatial, and other experiments suggest that these spatial codes are allocentric. PPA and RSC are more active when subjects are asked to report an object's location relative to an environment-centered reference frame than relative to either object- or viewer-centered reference frames (Committeri et al., 2004; Galati et al., 2010). In sum, PPA and RSC appear to be involved in encoding some aspect of space.

However, consistent with the neuropsychological findings, there is also evidence that PPA and RSC serve different functions and support different representations in the service of analyzing spatial information. In particular, PPA

appears to be more involved in coding aspects of the visible scene, whereas RSC appears to be more involved in coding the relationship of the scene to the broader environment. Although PPA is equally active when viewing any scene, RSC is significantly more active when subjects view familiar scenes for which they have long-term spatial knowledge (Epstein et al., 2007). Furthermore, although PPA is equally engaged by a variety of tasks involving scenes (Epstein et al., 2007), RSC is more strongly engaged by tasks that require encoding or retrieval of spatial information. For example, the amount of activity in RSC is correlated with the amount of survey knowledge subjects have learned while navigating a novel environment (Wolbers and Buchel, 2005), and activation in RSC is stronger when subjects are asked to retrieve allocentric spatial information about a familiar scene, such as location and heading, than when subjects are asked to make a familiarity judgment (Epstein et al., 2007). To better understand why PPA and RSC are differentially involved in these tasks, recent experiments have begun to probe the representational distinctions made by these regions.

There are two common techniques used to assess representations within a brain region: fMRI adaptation (fMRIa) and multivoxel pattern analysis (MVPA). Both techniques seek to determine which kinds of stimuli are coded as the “same” by a particular brain region and which kinds of stimuli are coded as “different.” In fMRIa, the assumption is that repetition of items that are representationally similar will lead to a decrease in neural response (Grill-Spector and Malach, 2001; Grill-Spector et al., 2006), though the mechanism by which this occurs is still a matter of debate (Sawamura et al., 2006; Summerfield et al., 2008; Larsson and Smith, 2012). In MVPA, the assumption is that similar items will elicit similar patterns of response across voxels; thus, this method makes use of multivariate information that is typically ignored in the standard univariate approach (Haxby et al., 2001; Norman et al., 2006; Mur et al., 2009).

Tests of representational distinctions in PPA suggest that it is sensitive to some spatial quantities of scenes, such as the spatial extent (Kravitz et al., 2011a; Park et al., 2014) or the precise viewpoint (Epstein et al., 2003; Epstein et al., 2008; Park and Chun, 2009), but is invariant to others, such as mirror-reversal (Dilks et al., 2011). PPA has also been implicated in coding of non-spatial aspects of scenes, such as texture (Cant and Xu, 2012) and category (Walther et al., 2009; Kravitz et al., 2011a). RSC is also sensitive to the spatial extent of a scene (Park et al., 2014) as well as its category (Walther et al., 2009), but in contrast to PPA, it is generally viewpoint-invariant (Epstein et al., 2008; Park and Chun, 2009) and sensitive to mirror-reversals of the image (Dilks et al., 2011). In sum, these results suggest that PPA and RSC support different representations during scene perception.

However, none of the previously described experiments tested whether PPA or RSC represent allocentric spatial information acquired from navigational experience. To date, only a handful of studies have investigated the representation of spatial information in the human brain, all of which have used virtual environments. Two studies have used fMRIa to test for heading-related codes. In the first study, subjects learned a virtual maze consisting of ten intersecting corridors, five of which were oriented north-south and five of which were oriented east-west (Baumann and Mattingley, 2010). At the end of each corridor was a landmark, which could only be viewed when facing a particular direction. When these landmarks were later shown in the scanner, a region in medial parietal cortex overlapping with RSC exhibited a reduction in response to consecutive landmarks that implied the same heading, indicating adaptation for facing direction. In the second study, subjects virtually navigated a circular arena with distal landmarks, collecting objects and replacing them based on their remembered locations (Doeller et al., 2010). They observed adaptation effects consistent with a population of direction-modulated grid cells: voxels in entorhinal cortex exhibited adaptation to running directions with six-fold rotational symmetry (e.g., if the initial running direction was 0°, adaptation was observed for runs at 60°, 120°, 180°, 240°, and 300°). Thus,

these findings are consistent with results from neurophysiological recordings, which have observed HD cells in retrosplenial cortex (Chen et al., 1994; Cho and Sharp, 2001) and grid cells in entorhinal cortex (Hafting et al., 2005; Jacobs et al., 2013).

Two studies have reported representations of spatial locations in the hippocampus (Hassabis et al., 2009; Rodriguez, 2010). In both of these studies, multivoxel patterns elicited during navigation in a small, simple environment were able to distinguish between 3 or 4 positions in the virtual room. However, an experiment using a more complex maze failed to observe multivoxel location coding in the hippocampus (Op de Beeck et al., 2013). Because the vast majority of place cell recordings have been obtained when animals explore small, simple environments (but see Kjelstrup et al., 2008), it is unknown whether place cells also support navigation on the very large spatial scales that humans typically traverse (e.g., a city). Alternatively, RSC might be better suited to support these representations, as is suggested by the neuropsychological literature.

We sought to address this gap by investigating which regions of the brain support representations of locations and headings within a very large environment for which subjects have extensive real-world navigation experience: their college campus. To this end, in three fMRI experiments, we scanned University of Pennsylvania students while they made decisions about places on campus. Because the Penn campus is laid out on a grid, location and heading can be easily determined at all times; as such, it is an ideal environment for testing for location and heading representations. We use both fMRIa and MVPA to test for these spatial representations, as both adaptation effects and multivoxel patterns can be simultaneously and independently measured when experiment trials are ordered according to a fully counterbalanced carryover sequence (Aguirre, 2007). In Chapter 2, we test for representations of locations (i.e., intersections) and allocentric headings during perception of visual scenes. In Chapter 3, we test whether these spatial representations are elicited during mental imagery and whether they take the same form as during perception. Finally, in Chapter 4, we test for metric

representations of location, i.e., representational similarity that scales with distance. Taken together, the work presented in Chapters 2-4 helps elucidate the role of PPA, RSC, the hippocampus, and other medial parietal and medial temporal regions in representing environmental space.

CHAPTER 2: ABSTRACT REPRESENTATIONS OF LOCATION AND FACING DIRECTION IN THE HUMAN BRAIN

Vass, L.K. and Epstein, R.A. (2013) Abstract representations of location and facing direction in the human brain. *Journal of Neuroscience*, 33(14): 6133-42.

2.1 Abstract

Humans, like other mobile organisms, rely on spatial representations to guide navigation from place to place. Although previous work has identified neural systems involved in wayfinding, the specific spatial codes supported by these systems are not well understood. We use functional magnetic resonance imaging (fMRI) to identify regions within the human medial temporal and medial parietal lobes that encode two fundamental spatial quantities – location and facing direction – in a manner that abstracts away from sensory inputs. University students were scanned while viewing photographs taken at several familiar campus locations. Multivoxel pattern analyses indicated that the left presubiculum, retrosplenial complex (RSC), and parietal-occipital sulcus (POS) coded location identity even across non-overlapping views, whereas the right presubiculum coded facing direction even across non-contiguous locations. The location and direction codes supported by these regions may be critical to our ability to navigate within the extended environment and to understand its large-scale spatial structure.

2.2 Introduction

Spatial navigation – getting from point A to point B in large-scale space—is a challenge that must be addressed by all mobile organisms. To solve this problem, many animals, including humans, rely on representations of the large-scale spatial structure of the world, mentally dividing their environment into different locations (or “places”) and remembering directional relationships between them. To use this

knowledge, a navigator must be able to represent its current location and facing direction and imagine other locations and facing directions. Here we investigate the neuroanatomical substrates of these location and direction codes.

Previous neurophysiological work, primarily in rodents, has implicated medial temporal lobe and Papez circuit structures in the coding of location and direction, revealing cell types that represent location ("place cells"; O'Keefe and Dostrovsky, 1971) facing direction ("head direction cells"; Taube et al., 1990b), and distances between points in space ("grid cells"; Hafting et al., 2005). The human neuroimaging (Epstein et al., 2008) and neuropsychological (Aguirre and D'Esposito, 1999) literatures, on the other hand, tend to implicate retrosplenial and medial parietal cortices (a region we label the "retrosplenial complex," or RSC) in addition to the medial temporal lobe. Both RSC and medial temporal lobe regions (hippocampus and parahippocampal cortex) activate during virtual navigation (Aguirre et al., 1996; Maguire et al., 1998; see Figure 2.1), and RSC is especially strongly engaged during retrieval of long-term spatial knowledge about familiar environments (Wolbers and Buchel, 2005; Epstein et al., 2007). Furthermore, damage to RSC leads to a profound inability to understand the spatial relationships between locations (Takahashi et al., 1997), a deficit that is not observed after damage to the hippocampus (Teng and Squire, 1999; Rosenbaum et al., 2000) or parahippocampal cortex (Habib and Sirigu, 1987). Thus, the neuropsychological and neuroimaging data (Maguire, 2001), together with studies of anatomical and functional connectivity (Kravitz et al., 2011b), support a model of human spatial navigation in which both medial temporal regions and RSC play central roles (Byrne et al., 2007; Vann et al., 2009). However, information about the specific spatial codes supported by these regions in humans has been sparse (although see Baumann and Mattingley, 2010).

Here we address this lacuna by using multivoxel pattern analysis (MVPA) of fMRI data to identify the regions that encode location and facing direction in the human brain. We scanned University of Pennsylvania students while they viewed photographs

taken facing different directions at several different locations around the Penn campus. We then attempted to decode location and facing direction based on multivoxel codes elicited while viewing these photographs. To anticipate, our results show that distributed activity patterns in RSC and the left presubicular region within the medial temporal lobe contain information about location whereas activity patterns in the right presubicular region contain information about facing direction; furthermore, they represent this information in a way that abstracts away from sensory features and thus is purely spatial.

2.3 Materials and Methods

SUBJECTS. Fourteen healthy subjects (9 female, mean age = 22 ± 0.5 y) with normal or corrected-to-normal vision were recruited from the University of Pennsylvania community and scanned with fMRI. An additional two subjects were scanned, but their data were not analyzed because of excessive head motion ($N=1$) or falling asleep during the experiment ($N=1$). All subjects were either current upper-class undergraduate students or recent graduates and all had extensive knowledge of the campus (mean years of campus experience = 2.9 ± 0.2). Prior to scanning, we confirmed their knowledge of the campus by asking them to indicate the location (e.g., 34th and Walnut St.) and facing direction (e.g., North) of photographs depicting 8 campus intersections not used in the main experiment. Subjects provided written informed consent in compliance with procedures approved by the University of Pennsylvania Institutional Review Board.

MRI ACQUISITION. Scanning was performed at the Hospital of the University of Pennsylvania using a 3T Siemens Trio scanner equipped with a 32-channel head coil. High-resolution T1-weighted anatomical images were acquired using a three-dimensional magnetization-prepared rapid-acquisition gradient echo pulse sequence [repetition time (TR) = 1620 ms; echo time (TE) = 3 ms; inversion time = 950 ms; voxel

size = 0.9766 x 0.9766 x 1 mm; matrix size = 192 x 256 x 160]. T2*-weighted images sensitive to blood oxygenation level-dependent contrasts were acquired using a gradient-echo echo-planar pulse sequence (TR = 4000 ms; TE = 30 ms; voxel size = 3 x 3 x 2 mm; matrix size = 64 x 64 x 65).

STIMULI AND PROCEDURE. Stimuli consisted of digitized color photographs taken facing the 4 compass directions (North, East, South, West) at 8 intersections on the Penn campus (Figure 2.2). For each of these 32 views (8 locations x 4 directions), 17 different photographs were acquired, making a total of 544 images in all. The stimulus set was split in half such that subjects saw all 272 images corresponding to 4 of the intersections (from Figure 2.2, intersections 1, 3, 6, and 8) twice over the course of 4 runs lasting 10 min 52 s each and then saw all 272 images corresponding to the 4 remaining intersections (from Figure 2.2, intersections 2, 4, 5, and 7) twice during 4 more runs of equal length.

Each scan run of the main experiment was divided into 137 or 139 stimulus trials interspersed with 7 or 8 null trials. On each stimulus trial, subjects viewed a single image from the stimulus set at 1024 x 768 pixel resolution for 3 s followed by a 1 s interstimulus interval. Their task was to determine the direction that the camera was facing when the image was taken (North, East, South, or West). Because subjects typically find it impossible to identify facing direction without first identifying the location of the image, performance of this task ensures that neural systems representing both location and facing direction are activated. Subjects responded either by pressing the button on a 4-button response box corresponding to that direction or by covertly identifying the direction and making a button press (i.e., the same button for all directions). Null trials consisted of 8 s of a gray screen with black fixation cross during which subjects made no response.

Trials were ordered according to a continuous carryover sequence—a serially balanced design in which each view follows every other view including itself exactly

once (Aguirre, 2007). Subjects viewed 2 full carryover sequences for each set of intersections, with each sequence spanning 2 scan runs. A unique carryover sequence was generated for each subject for each set of intersections. Subjects switched response instructions after every carryover sequence (i.e., every 2 scan runs) so that facing direction was not confounded by motor response in the multivoxel pattern analyses.

In addition to the main experiment, subjects also completed a functional localizer scan lasting 9 min 52 s which consisted of 16-s blocks of scenes, objects, and scrambled objects. Images were presented for 490 ms with a 490 ms ISI as subjects performed a one-back task on image repetition.

DATA PREPROCESSING. Functional images were corrected for differences in slice timing by resampling slices in time to match the first slice of each volume. Images were then realigned to the first volume of the scan run and spatially normalized to the Montreal Neurological Institute template. Data for the functional localizer scan were smoothed with a 6-mm full-width-half-maximum Gaussian filter; data for multivoxel pattern analyses were not smoothed.

FUNCTIONAL REGIONS OF INTEREST. Data from the functional localizer scan were used to identify two scene-responsive regions, the parahippocampal place area (PPA) and retrosplenial complex (RSC), which have been previously implicated in place recognition and navigation. These regions were defined as the set of contiguous voxels in the vicinity of the parahippocampal/lingual boundary (PPA) or retrosplenial cortex/parietal-occipital sulcus (RSC) that responded more strongly to scenes than to objects. Thresholds were determined separately for each subject to be consistent with ROIs found in previous studies and ranged from $t > 1.5$ to $t > 3.5$ (mean $t = 2.7 \pm 0.2$). Data from one subject's functional localizer scan could not be used, so for this subject we used the across-subject ROI intersection that most closely matched the average size of

each ROI in the remaining 13 subjects. We identified bilateral PPA in all 13 of the remaining subjects, left RSC in 11/13 and right RSC in 12/13.

ANATOMICAL REGIONS OF INTEREST. Because neurons sensitive to spatial quantities have been identified throughout the medial temporal lobe, we created six anatomical ROIs that covered this region: anterior hippocampus, posterior hippocampus, presubiculum (a structure on the medial parahippocampal gyrus situated between the subiculum and entorhinal cortex), entorhinal cortex, perirhinal cortex, and parahippocampal cortex (Figure 2.4a). The hippocampus and presubiculum were defined using the fully automated segmentation protocol in FreeSurfer 5.1 (Van Leemput et al., 2009). This technique uses Bayesian inference on an upsampled version of the T1 structural image to determine the likely hippocampal subfield identity of $0.5 \times 0.5 \times 0.5$ mm voxels. We first assigned a subfield identity to each “mini-voxel” by selecting the subfield with the highest probability. We then assigned the identity of each $3 \times 3 \times 2$ mm functional voxel according to the most commonly occurring subfield across the 144 mini-voxels. The hippocampus ROI was defined as the union of the CA1, CA2/3, CA4/Dentate Gyrus, and subiculum. We then divided the hippocampus into anterior and posterior subregions at $y = -26$. Entorhinal, perirhinal, and parahippocampal cortices were defined based on manual parcellation of the T1 anatomical image following the protocol in Pruessner et al. (2002), with the additional constraint that the parahippocampal cortex ROI did not contain any PPA voxels.

MULTIVOXEL PATTERN ANALYSIS. To determine whether each ROI contained information about location, facing direction, or the identity of the specific view, we calculated correlations between multivoxel patterns elicited in different carryover sequences (Haxby et al., 2001). Data from the first half of the experiment (in which 4 of the 8 intersections were shown in two carryover sequences) were analyzed separately

from data from the second half of the experiment and the results were averaged together. Below we describe the analysis procedure for each half.

The first step of the analysis was to obtain the multivoxel activity patterns reflecting the response to each view for each carryover sequence. To create these, we first passed the timecourse of MR activity for each voxel through a general linear model (GLM) containing 32 regressors (16 views x 2 carryover sequences). The resulting β values corresponded to an estimate of the average response to each view within each carryover sequence. Multivoxel activity patterns were then constructed by simply concatenating these response values across voxels. GLMs were implemented in VoxBo (www.voxbo.org) and included an empirically derived $1/f$ noise model, filters that removed high and low temporal frequencies, and nuisance regressors to account for global signal variations and differences in mean signal between scan runs.

The second step of the analysis was to assess similarity between activity patterns by calculating correlations between patterns in different carryover sequences (Figure 2.3). First, we normalized the patterns by calculating the mean pattern across all views separately for each carryover sequence and subtracting this mean pattern from each of the individual patterns. Next, we created a 16 x 16 correlation matrix by calculating the cross-sequence Pearson correlations between patterns for all pairs of views, including both same-view pairs (e.g., View 1 in Sequence 1 vs. View 1 in Sequence 2) and different-view pairs (e.g., View 1 in Sequence 1 vs. View 2 in Sequence 2). Each cell of the correlation matrix belonged to 1 of 4 possible groups based on shared spatial quantities: 1) Same Location Same Direction (SLSD; i.e., same view); 2) Same Location Different Direction (SLDD; e.g., Location 1 facing North vs. Location 1 facing East); 3) Different Location Same Direction (DLSD; e.g., Location 1 facing North vs. Location 2 Facing North); 4) Different Location Different Direction (DLDD; e.g., Location 1 facing North vs. Location 2 facing West).

The third step of the analysis was to evaluate whether the correlation values determined in step 2 reflected coding of view, location, or direction. To assess this, we

performed 6 separate multiple regressions on these correlation values. Three of these models examined view, location, and direction coding before controlling for visual similarity; the remaining 3 examined view, location, and direction coding after controlling for visual similarity. The independent variables in the first set of models (before controlling for visual similarity) were a binary categorical regressor for the contrast of interest and a constant term; the models that controlled for visual similarity contained the same regressors plus a parametric regressor for visual similarity (see below for details). To test for coding of view, we used a categorical regressor that contrasted between SLSD and the average of SLDD, DLSD, DLDD—that is, same view versus all of the different-view groups. To test for coding of location independent of view, we used a categorical regressor that contrasted SLDD with DLDD—that is, same versus different location under the constraint that direction (and hence view) is always different. To test for coding of direction independent of view, we used a categorical regressor that contrasted DLSD with DLDD—that is, same vs. different direction under the constraint that location (and hence view) is always different. These analyses were performed on the full dataset for 11 of the 14 subjects; for the remaining 3 subjects, they were performed on a partial dataset consisting of only one half of the experiment, because these subjects either did not complete both halves of the experiment (N=1) or fell asleep during scan runs from one half of the experiment (N=2). All independent variables were standardized before running the model by subtracting the mean of the regressor and then dividing by the standard deviation of the regressor.

To create a parametric regressor for visual similarity, we calculated the visual similarity between all pairs of images using a simple texture model that has previously been shown to approximate human performance on scene identification at very brief (<70 ms) exposures (Renninger and Malik, 2004). Images were first converted to grayscale and then passed through V1-like Gabor filters of varying orientations and sizes in order to identify the 100 most common texture features across images (Matlab code available at renningerlab.org). For each image, we generated a histogram

reflecting the frequency of each of the 100 texture features. For every pair of images, we computed visual dissimilarity by comparing the distributions of texture features using a χ^2 measure (smaller χ^2 corresponds to less visual dissimilarity). Then, to calculate visual dissimilarity between a pair of views, we averaged over all the relevant pairwise combinations of images. Finally, we converted the visual dissimilarity values to visual similarity by subtracting each χ^2 from the maximum χ^2 .

SEARCHLIGHT ANALYSIS. To test for coding of spatial quantities outside of our ROIs, we implemented a searchlight procedure (Kriegeskorte et al., 2006), which performs the same calculations described above, but in small spherical ROIs (radius=5 mm) centered on each voxel of the brain in turn. Thus, we determined the average correlation between each pair of views in the local neighborhood of each voxel, performed the multiple regressions as described above, and assigned the beta weight for the regressor of interest to the center voxel of the spherical ROI. This procedure generated 6 whole-brain maps for each subject corresponding to coding for view, location, and direction in each half of the experiment. For each type of spatial coding, we averaged together the maps from each half of the experiment and then submitted this average map to a second-level random-effects analysis to identify voxels that reliably exhibit spatial coding across subjects. To find the true Type I error rate for each type of spatial coding, we performed Monte Carlo simulations, which involved sign permutations of the whole-brain data from individual subjects (Nichols and Holmes, 2002). We then report voxels which are significant at $P < 0.05$ after correcting for multiple comparisons across the entire brain.

TUNING FOR FACING DIRECTION. To further explore direction coding revealed by the main MVPA analyses described above, we calculated tuning curves for each direction to test whether all directions are equally-well represented. This involved a modification to steps 2 and 3 of the analysis procedure described above. After computing the 16 x

16 matrix for similarity between views, we grouped these correlation values according to their direction-direction pairing (e.g., East-East, East-North, etc.). Because we wanted to examine direction coding independent of location or view, we excluded same-location pairings (i.e., SLSD and SLDD) from these groupings. We then performed a multiple regression on the correlation values to obtain estimates of the average correlation for each of the ten direction-direction groupings while taking visual similarity into account. These data were then further analyzed by comparing the within-direction (DLSD) beta weight for each direction (North, South, East, West) to the between-direction betas (DLDD) for that direction (e.g., East-East vs. East-North, East-South, and East-West). Finally, we performed a 4 x 2 repeated measures ANOVA on direction (North, East, South, West) and pattern similarity type (Within Direction, Between Directions).

TUNING FOR LOCATION. To further explore location coding revealed by the main MVPA analyses describe above, we looked for evidence of graded coding of location by testing whether pattern similarity varied based on the real-world Euclidean distance between intersections. In contrast to the previous analyses, where we analyzed correlations separately for each set of 4 intersections and then averaged the results together, here we examined correlations between all 8 intersections in order to maximize the variability in between-location distances. We therefore excluded from this analysis 3 subjects for whom we did not have data from both halves of the experiment. For the remaining 11 subjects, we calculated the across-carryover-sequence correlations between all pairs of DLSD and DLDD views and then grouped together all view pairs from the same pair of locations. For example, the location pair “Location 1—Location 2” would include view pairs such as “Location 1 East—Location 2 North” and “Location 1 South—Location 2 West.” We then defined the neural distance between two locations as the average pattern dissimilarity (i.e., $1-r$) across all view pairs from that location pair. We then fit these neural distance values to a multiple regression model

that included regressors for visual similarity, the real-world Euclidean distance between locations, and a constant term. We extracted the beta weights for the Euclidean distance regressor for each subject and compared them to zero using a two-tailed t -test.

FMRI ADAPTATION. To test for a reduction in response after repetitions of view, location, or direction, we combined the data from all 8 intersections and created a model in which each trial is defined based on shared spatial quantities with the previous trial. Thus, there were regressors for Repeat View, Repeat Location (but change direction), Repeat Direction (but change location), and Nonrepeat trials (i.e., change direction and change location). The model also included a regressor for trials that followed null trials, a regressor that modeled low-level visual similarity, and nuisance regressors as described above. We extracted β values for each regressor and performed three planned two-tailed t tests. To test for effects of view repetition, we compared Repeat View to the average of Repeat Location, Repeat Direction, and Nonrepeat, since those three regressors all reflect activity during trials for which the view is different from the previous trial. To test for effects of location repetition, we compared Repeat Location to Nonrepeat. Finally, to test for effects of direction repetition, we compared Repeat Direction to Nonrepeat.

2.4 Results

BEHAVIORAL PERFORMANCE. On each trial, subjects reported the facing direction (North, South, East, West) for an image of the Penn campus. This task requires subjects to retrieve spatial information about the depicted location that goes beyond simple perceptual analysis of the image. Subjects performed this task rapidly and accurately (average reaction time (RT) for all trials = 1331 ± 62 ms; average accuracy for the 50% of trials in which subjects explicitly reported direction = $88.1 \pm 2.5\%$), which was expected given the grid-like structure of the Penn campus (Figure 2.2) and the fact that

we pre-screened subjects to ensure they could perform the task using 8 different campus intersections not shown in the main experiment.

Using these behavioral responses, we looked for evidence of behavioral priming by sorting the trials into four trial types based on whether the image shown on the trial depicted the same or different spatial information as the image shown on the immediately-preceding trial. On *Repeated View* trials, both location and direction were maintained across the successive trials (e.g., Intersection 1 Facing East → Intersection 1 Facing East). On *Repeated Location* trials, location was maintained, but direction differed (e.g., Intersection 1 Facing East → Intersection 1 Facing North). On *Repeated Direction* trials, direction was maintained, while location differed (e.g., Intersection 1 Facing East → Intersection 2 Facing East). Finally, on *Nonrepeat* trials, both location and direction changed (e.g., Intersection 1 Facing East → Intersection 2 Facing South). We then compared the average reaction times across the four trial types to look for evidence of priming for view, direction, or location. Note that the images shown on successive trials were never exactly identical, even on repeated view trials, as each individual image was shown only twice during the experiment and never more than once within the same scan run.

This analysis revealed behavioral priming effects for repetition of the same view (Repeated View RT, 1218 + 56 ms vs. Change View RT [average of Repeated Location, Repeated Direction, and Nonrepeat RTs], 1368 + 64 ms; $t_{13} = -7.37$, $P = 0.000005$) and repetition of the same location (Repeated Location RT, 1346 + 66 ms vs. Nonrepeat RT, 1377 + 66 ms; $t_{13} = -3.2$, $P = 0.007$), but no priming for repetition of the same direction (Repeated Direction RT, 1383 + 61 ms vs. Nonrepeat RT, 1377 + 66 ms; $t_{13} = 0.6$, $P = 0.54$). Thus, subjects exhibited priming when either view or location was repeated on successive trials, even though they were performing a direction task. The presence of a location priming effect and absence of a direction priming effect is not surprising given that subjects typically find it impossible to identify facing direction without first identifying location. Thus, when location is repeated across trials, subjects exhibit a

benefit because they accessed the same location on the previous trial. However, when direction is repeated, it provides no benefit since direction cannot be directly ascertained from the image.

MULTIVOXEL DECODING OF VIEW, LOCATION, AND FACING DIRECTION. We then turned to the main question of the study: how is information about views, locations, and facing directions encoded in different brain regions? To assess this, we measured similarity between multivoxel activity patterns evoked by different stimuli. We hypothesized that if a region contains information about a particular spatial quantity, such as location, the evoked activity patterns for two stimuli that share that quantity (e.g., *same location*) should be more similar than the evoked activity patterns for two stimuli that do not share that quantity (e.g., *different locations*). Because previous work suggested that the parahippocampal place area (PPA) and retrosplenial complex (RSC) might be especially involved in coding spatial quantities, we first focused on these regions before examining medial temporal lobe regions, and finally considering patterns in all regions of the brain.

We first performed three multiple regressions on the pattern similarity data to test for coding of views, locations, and directions, respectively (Figure 2.3). Patterns in PPA distinguished between views ($t_{13}=5.8$, $P<0.0001$) and locations ($t_{13}=2.3$, $P=0.04$), but distinctions between directions ($t_{13}=2.0$, $P=0.07$) fell short of significance. That is, patterns elicited by the same view in different scan runs were more similar than patterns elicited by different views, and patterns elicited by different views of the same location were more similar than patterns elicited by different locations. Patterns in RSC not only distinguished between specific views ($t_{12}=5.7$, $P<0.0001$) and locations ($t_{12}=4.3$, $P=0.001$), but also distinguished between facing directions ($t_{12}=2.8$, $P=0.015$). Because the multivoxel patterns compared in the same view condition reflect the response to the same set of images, view coding could not be fully disentangled from image coding or visual feature coding in this experiment. In contrast, the analyses for

location and direction coding were designed to minimize the effects of visual similarity since the pairs of views in the *same location* (SLDD) and *same direction* (DLSD) conditions did not contain overlapping visual information. However, this design may not have fully controlled for effects of visual similarity because there might be low level visual features in common across same-location and same-direction images that could give rise to the increased pattern similarity. To account for this possibility, we performed three additional multiple regressions for view, location, and direction, each of which included a covariate for low-level visual similarity (see Materials and Methods for more details on the visual similarity measure). After controlling for visual similarity, activity patterns in PPA distinguished between views ($t_{13}=5.1$, $P=0.0002$), but did not distinguish between locations ($t_{13}=1.5$, $P=0.15$) or directions ($t_{13}=0.4$, $P=0.71$; Figure 2.3). Activity patterns in RSC distinguished between views ($t_{12}=3.8$, $P=0.003$) and locations ($t_{12}=3.3$, $P=0.006$), but the distinction between directions was reduced to a marginal trend ($t_{12}=1.9$, $P=0.09$). Thus, once visual similarity is controlled for, PPA no longer exhibits coding for spatial quantities, and the evidence for coding of direction in RSC becomes less clear. Because visual similarity affected our estimates of pattern similarity in PPA and RSC, we included visual similarity as a covariate in all subsequent analyses of pattern similarity.

We then looked for evidence of spatial coding in medial temporal lobe regions. Rodent extracellular recordings strongly implicate these regions in coding for location (hippocampus; O'Keefe and Dostrovsky, 1971) and head direction (presubiculum; Taube et al., 1990b). However, the extent to which these regions are involved in spatial coding in humans is unclear. The hippocampus is activated in some neuroimaging studies of navigation (Ghaem et al., 1997; Maguire et al., 1998), but not in others (Aguirre et al., 1996; Aguirre and D'Esposito, 1997; Rosenbaum et al., 2004). Furthermore, although putative place cells have been identified in this region (Ekstrom et al., 2003), humans with hippocampal damage retain the ability to navigate through familiar environments (Teng and Squire, 1999). As for human presubiculum, there is

currently no evidence implicating this region in coding for head direction, though HD cells have been identified there in nonhuman primates (Robertson et al., 1999). Because the contributions of these and nearby medial temporal regions to human spatial coding are yet unresolved, we specifically targeted them in the following analyses.

Anterior and posterior hippocampus (including the CA fields, dentate gyrus, and subiculum) and extra-hippocampal regions (presubiculum, entorhinal cortex, perirhinal cortex, and parahippocampal cortex exclusive of PPA) were defined anatomically as described in Materials and Methods (Figure 2.4). We then performed the same tests for view, location, and direction coding described above, and submitted the results of each test to one-way repeated measures ANOVAs with Region as a factor. These analyses revealed that coding for location and direction did indeed differ between ROIs, as evidenced by significant effects of Region (Location: $F_{5,65}=4.3$, $P=0.002$; Direction: $F_{5,65}=4.7$, $P=0.001$); coding for view did not differ between regions ($F_{5,65}=0.5$, $P=0.78$). When we explicitly tested the location, direction, and view effects within each region, we found that presubiculum coded for direction ($t_{13}=3.0$, $P=0.01$) and location ($t_{13}=3.3$, $P=0.006$). No other regions coded for direction or location, nor did any region distinguish between different views (Figure 2.4).

These results are consistent with animal neurophysiology literature indicating that presubiculum contains a mixture of cells that convey information about direction and location, including HD cells (Boccarda et al., 2010), grid cells (ibid.), and theta-modulated place-by-direction cells (Cacucci et al., 2004). Inspection of the data suggested that presubiculum coding for these two types of information differed across hemispheres. To further characterize these hemispheric effects, we performed a 2 x 2 repeated measures ANOVA with Hemisphere and Information Type (location and direction) as factors. There were no significant main effects, but there was a significant interaction between Hemisphere and Information Type ($F_{1,13}=7.5$, $P=0.02$), indicating that the degree to which presubiculum coded for location and direction differed by

hemisphere. Specifically, left presubiculum distinguished between locations (t -test vs. 0, $t_{13}=3.3$, $P=0.005$) but not directions ($t_{13}=1.7$, $P=0.11$) whereas right presubiculum distinguished between directions ($t_{13}=3.4$, $P=0.005$) but not locations ($t<1$, n.s.; Figure 2.4). Furthermore, direction coding was stronger than location coding in right presubiculum ($t_{13}=2.3$, $P=0.04$), although the reverse was not true for left presubiculum ($t_{13}=1.7$, $P=0.11$).

WHOLE-BRAIN SEARCHLIGHT ANALYSES. Having identified multivoxel activity patterns corresponding to coding of view, location, and direction in our pre-selected ROIs, we then performed a searchlight analysis (Kriegeskorte et al., 2006) to determine whether these quantities could be decoded elsewhere in the brain. The resulting significance maps confirmed the effects found in the ROI analysis and identified additional areas that show sensitivity to these quantities (Figure 2.5).

Location could be decoded from multivoxel patterns within a swath of cortex along the parietal-occipital sulcus (POS) that overlapped with posterior RSC and continued posteriorly and superiorly into the precuneus. In contrast, direction was decodable in more anterior regions that included the anterior calcarine sulcus and a region partially overlapping with the posterior presubiculum. Finally, individual views could be decoded from multivoxel patterns throughout visual cortex including early visual cortex and territory in the object-selective lateral occipital complex. The fact that views could be decoded in early visual regions suggests that there are similarities between same-view images that are not captured by our visual similarity model, which focuses on texture similarities without consideration of color or the spatial distribution of features within the image.

TUNING FOR FACING DIRECTION. To further investigate the nature of direction coding in right presubiculum we constructed directional tuning curves for this region (Figure 2.6). This involved plotting the correlations for each direction pairing separately, rather

than averaging over all same-direction and all different-direction view pairings as we did previously. As in the earlier analyses, we considered correlations between views from different locations to ensure that similarities in direction were not confounded by similarities in view or location, and we included a regressor to control for low-level visual similarity.

To test whether right presubiculum exhibited directional tuning for all directions or just a subset of directions, we performed a 4 (North, East, South, West) x 2 (Within-Direction, Between-Directions) repeated measures ANOVA comparing pattern similarity within a direction (e.g., East-East) to pattern similarity between directions (e.g., East-North, East-South, East-West), separately for each direction. Pattern similarities showed a main effect of Within- vs. Between-Directions ($F_{1,13}=10.6$, $P=0.006$), confirming that this region distinguishes between directions. There was no interaction effect ($F_{3,39}=0.6$, $P=0.61$), suggesting that right presubiculum does not show preferential tuning for a subset of the 4 directions. Thus, the directional tuning observed here is similar to that observed in rodent HD cells insofar as all directions are equally-well represented.

TUNING FOR LOCATION. The location coding results reported above indicate that the multivoxel activity patterns in RSC, left presubiculum, and POS are capable of distinguishing between different campus locations. Yet this analysis does not reveal whether these locations are represented in a map-like (i.e. metric) fashion, whereby locations that are closer together in the real world are representationally more similar than locations that are further apart. We have previously shown using fMRI adaptation that the bulk activity in left anterior hippocampus is parametrically modulated by the real-world distance between places (Morgan et al., 2011), but this type of coding has not yet been demonstrated using multivoxel patterns. To test this possibility, we regressed the “neural distance” between activity patterns against the real-world Euclidean distance between the corresponding locations. Neural distance was defined as the average pattern dissimilarity (i.e., $1-r$) between all pairs of views belonging to a

pair of locations. As in the previous analyses, this model also included a regressor for the low-level visual similarity between locations. We performed this analysis in regions that demonstrated a location effect in the earlier analyses, but found no evidence for this type of coding in RSC (mean beta = 0.05, $P=0.10$), left presubiculum (mean beta = -0.004, $P=0.79$), or the region along the POS identified in the location searchlight analysis (mean beta = 0.04, $P=0.16$). These results suggest that although distributed patterns in these regions can distinguish between real-world locations, they might not contain information about the underlying metric structure of the environment. When we repeated this analysis across the entire brain using a searchlight method, we did not find any voxels that demonstrated a relationship between pattern similarity and Euclidean distance between locations. Nor did we find any effects in PPA or the remaining medial temporal lobe ROIs (all $P_s > 0.32$).

FMRI ADAPTATION. A second technique used to infer representational distinctions within a region is fMRI adaptation—a reduction in response when stimulus features are repeated (Grill-Spector et al., 2006). Previous studies that have assessed neural representations using both MVPA and fMRI adaptation have reported inconsistencies in the results obtained by these two techniques (Drucker and Aguirre, 2009; Epstein and Morgan, 2012). These discrepancies do not appear to reflect a difference in sensitivity between MVPA and fMRI adaptation (but see Sapountzis et al., 2010), but rather suggest a difference in the precise features of the neuronal code interrogated by these two techniques (Drucker and Aguirre, 2009; Epstein and Morgan, 2012). As such, fMRI adaptation analyses have the potential to provide complementary results to those obtained by MVPA.

Because we used a continuous carryover sequence in which each view was presented after every other view including itself equally often, adaptation effects were independent of the main effects analyzed by MVPA in our experiment and could be independently assessed. We therefore examined the effects of repeating view, location,

or direction across pairs of trials. As in the analysis of the behavioral data, we sorted trials into 4 trial types based on the shared spatial information with the immediately-preceding trial: *Repeated View*, *Repeated Location*, *Repeated Direction*, or *Nonrepeat*. We then looked for a reduction in fMRI response in the current trial caused by repetition of spatial information from the previous trial.

We first looked for adaptation to view repetitions by comparing the mean fMRI signal of *Repeated View* trials to the mean fMRI signal of all different view trials (i.e., the average of *Repeated Location*, *Repeated Direction*, and *Nonrepeat*). In line with previous experiments, we found robust adaptation to repeated view in PPA and RSC (PPA: $t_{13}=-9.5$, $P<0.0001$; RSC: $t_{12}=-6.6$, $P<0.0001$). In addition, many of the medial temporal lobe ROIs, with the exception of entorhinal and perirhinal cortices, also exhibited adaptation to repeated view (Anterior Hippocampus: $t_{13}=-2.5$, $P=0.03$; Posterior Hippocampus: $t_{13}=-4.5$, $P=0.0006$; Presubiculum: $t_{13}=-4.9$, $P=0.0003$; Parahippocampal Cortex: $t_{13}=-5.1$, $P=0.0002$). We next looked for adaptation to repetitions of location (under the constraint that direction always differed) by comparing the *Repeated Location* trials to the baseline *Nonrepeat* trials. Although the response to *Repeated Location* was numerically smaller than response to *Nonrepeat* trials in RSC, there were no significant effects of repeating location in this or any other ROI (all $P_s>0.11$).

Finally, we looked for adaptation to repetitions of direction (under the constraint that location always differed) by comparing the *Repeated Direction* trials to the baseline *Nonrepeat* trials. Here, we find an unexpected effect of anti-adaptation (i.e., greater response when direction was repeated) in RSC ($t_{12}=2.9$, $P=0.01$) and parahippocampal cortex ($t_{13}=2.2$, $P=0.048$), and a trend for anti-adaptation in PPA ($t_{13}=1.8$, $P=0.09$). We observed no effects of repeating facing direction in the remaining ROIs (all $P_s>0.56$). A recent fMRI adaptation experiment observed reduced response in a medial parietal region within RSC when facing direction was repeated across trials—the opposite effect to what we observe here (Baumann and Mattingley, 2010). An important difference

between the two studies is that subjects in the previous experiment performed an orthogonal location task, whereas here subjects actively reported facing direction on every trial. Because our task led subjects to focus on direction, the repeated direction trials may have been particularly salient and engaged additional attention and hence increased rather than reduced response, a hypothesis supported by predictive coding models of anti-adaptation (Segaert et al., 2013).

In sum, as observed in previous studies, fMRI adaptation results were partially consistent and partially inconsistent with those obtained by MVPA. Both methods found evidence for view coding in PPA and RSC and some degree of evidence for direction coding in RSC. However, location coding in RSC and medial temporal lobe regions was only observed with MVPA while view coding in medial temporal lobe ROIs was only observed with adaptation. These findings are broadly consistent with previous results indicating that fMRI adaptation effects tend to index representational distinctions that are more stimulus-specific than those indexed by MVPA, either because the locus of the effect is closer to the single unit (Drucker and Aguirre, 2009), or even the synapse (Sawamura et al., 2006), or because it reflects the operation of dynamic recognition mechanisms that are more tied to individual items than to general categories (Summerfield et al., 2008; see Epstein and Morgan, 2012 for additional discussion).

2.5 Discussion

The principal finding of this study is that distributed patterns of fMRI activity in the medial temporal lobe and medial parietal cortex contain information about location and facing direction within a familiar, real-world environment. Specifically, patterns in RSC, POS, and left presubiculum contain information about location while patterns in right presubiculum contain information about facing direction. These results demonstrate the coding of navigationally-relevant spatial information in specific regions in the human brain.

Our first main finding is that locations are represented in both medial temporal (left presubiculum) and medial parietal regions (RSC and more posterior territory along the POS, i.e., precuneus). This was demonstrated by the finding of similarity between multivoxel patterns elicited by different views taken at the same intersection. This result cannot be explained by coding of visual features since we compared pairs of views with non-overlapping visual information and explicitly modeled low-level visual similarity between the views. Indeed, it would be virtually impossible for an observer to know which views corresponded to which locations without long-term knowledge about the spatial location of the views. Nor can the results be explained by coding of behavioral responses corresponding to the different locations because the task required subjects to report facing direction, not location. Rather, these findings demonstrate the existence of a purely abstract representation of location in medial temporal and medial parietal regions whereby different topographical features corresponding to the same location elicit a common neural response.

By showing a neural representation of location per se in RSC and POS, these findings go beyond previous multivoxel studies that have demonstrated decoding of scene categories (Walther et al., 2009), the geometric structure of scenes (Kravitz et al., 2011a; Park et al., 2011) and individual landmarks (Morgan et al., 2011; Epstein and Morgan, 2012) in several cortical regions, including the PPA and RSC. Furthermore, whereas previous neuroimaging work has implicated RSC in the retrieval of spatial quantities—for example by showing that fMRI activity in RSC increases when subjects retrieve long-term spatial knowledge (Epstein et al., 2007) or move through an environment for which they have obtained survey knowledge (Wolbers and Buchel, 2005)—those studies could not exclude the possibility that RSC mediates general mnemonic processes that facilitate encoding or retrieval. The current findings provide more convincing evidence for the proposition that RSC represents location, since multivoxel codes in RSC (and POS) reliably distinguished between locations, even though location was not reducible to differences in visual features. In contrast,

multivoxel codes in the PPA distinguished between views but did not distinguish between locations when visual similarity was controlled. These findings are consistent with the view that PPA is a visual region that primarily represents the local scene, whereas RSC and other parietal regions support spatial representations that connect the local scene to the broader environment (Epstein et al., 2007; Park and Chun, 2009). The finding of location coding in RSC/POS fits well with a previous report of location-sensitive neurons in monkey medial parietal cortex (Sato et al., 2006). In that study, monkeys navigated along different routes to goal locations in a virtual environment, and a subset of medial parietal neurons fired whenever the monkey occupied a particular virtual location, regardless of the route. Compared to place cells in the hippocampus, these place-related cells in the medial parietal lobe have been less studied, and less is known about the precise spatial quantities they might encode. One possibility is that location representations in medial parietal cortex might be more schematic whereas location representations in hippocampus might be more metric. Consistent with this view, similarities between location patterns in RSC and POS did not relate to real-world Euclidean distances between locations. Although we did not detect location coding in the hippocampus in the current experiment, we have previously shown using an adaptation paradigm that the hippocampus has access to metric information about which locations are closer to each other and which are further away (Morgan et al., 2011). Thus, RSC and POS might encode locations (and, possibly, directional vectors between locations) but might not organize locations according to their exact coordinates within a continuous spatial map (Gallistel, 1990). This could explain previous findings that hippocampal amnesic patients (whose RSC is intact) can navigate premorbidly familiar environments, but display deficits when the navigation task requires fine-grained spatial information (Teng and Squire, 1999; Maguire et al., 2006). In any case, the current observation of abstract coding of location in RSC and POS indicates that place representations are not restricted to the medial temporal lobe, but can also be found in medial parietal regions.

We also observed location coding in presubiculum, a finding that is consistent with results from rodent neurophysiology studies, which have identified cells that exhibit location coding in this region, including grid cells (Boccaro et al., 2010) and theta-modulated place-by-direction cells (Cacucci et al., 2004). This result resembles a previously-reported finding that multivoxel patterns in the hippocampus could distinguish between corners of a newly-learned virtual room (Hassabis et al., 2009). Complementing this, here we show that fMRI patterns in left presubiculum (in addition to RSC/POS) distinguish between real-world locations within a much larger environment (350 x 270 m in the current study versus 15 x 15 m in the previous study) for which the subjects have years of experience. Differences in size and familiarity between the environments might partially explain why spatial codes were found in different regions in the two experiments: the hippocampus may be more important for coding location within small and/or newly-learned environments, whereas presubiculum and RSC/POS may be more important for coding location within large and/or highly familiar environments (see Smith et al., 2012). Alternatively, the apparent discrepancies between experiments might be attributable to differences in ROI definition and fMRI acquisition parameters: the previous study did not attempt to distinguish between coding in the presubiculum and coding in the hippocampus and did not acquire data from RSC; on the other hand, it used smaller voxels that might have made it more sensitive to hippocampal differences. Further experiments are needed to fully characterize the properties of the location codes in these regions and under what circumstances they arise.

The second main result of the current study is that allocentric facing directions are represented in the right presubiculum. This was demonstrated by the increased pattern similarity between views facing the same direction (e.g., North) across different intersections. Although this finding is somewhat confounded by the fact that our subjects were performing a direction task during the experiment, it cannot be simply explained in terms of response preparation or execution because we always compared

neural activity across the two response modalities—subjects either indicated facing direction via button press or covertly identified the facing direction. These results fit well with the animal literature since HD cells have been found in the presubiculum of rats (Taube et al., 1990b) and nonhuman primates (Robertson et al., 1999). Furthermore, the presubiculum is thought to be critical for updating the HD signal based on visual landmark information (Yoder et al., 2011), which was the only directional cue available in this experiment, given the absence of vestibular and self-motion information in the scanner.

Some of our MVPA data suggested that RSC might also code for allocentric facing direction; however, this result did not maintain significance when low level visual similarity between directions was controlled. The human literature supports this role for RSC since patients with damage to retrosplenial cortex are described as having "lost their sense of direction" (Takahashi et al., 1997) and medial parietal cortex exhibits neural adaptation to repetitions of heading directions (Baumann and Mattingley, 2010). Indeed, we also find differences in RSC activity after direction is repeated, though we observed an increase (i.e., anti-adaptation) rather than a decrease in activity. There is also evidence from rodent studies for HD cells in RSC (Chen et al., 1994; Cho and Sharp, 2001). Thus, the failure to observe significant distributed direction coding in RSC should not be taken as evidence that this region does not represent this information. Rather, it may simply reflect the fact that allocentric direction was correlated with low level visual properties in our stimulus set. Alternatively, it is possible that RSC might have coded direction separately for each intersection rather than using a single representation that applied across the entire campus, or coded direction in terms of distal landmarks which were only partially consistent across locations. These results warrant future experiments using different stimuli for which allocentric direction is fully unconfounded from these factors.

In summary, our study demonstrates neural coding of spatial quantities necessary for navigation within a large-scale, real-world environment. We find that

location and facing direction are represented in distributed patterns in specific brain regions, with medial temporal and medial parietal regions coding location, and medial temporal regions coding facing direction. These results provide an important link between the animal neurophysiology, human neuropsychology, and human neuroimaging literatures, and are a first step toward understanding how spatial information might be represented in both the medial temporal lobe and higher-level association areas.




Figure 2.1 Meta-analysis of brain regions activated during studies of navigation. Neurosynth (www.neurosynth.org; Yarkoni et al., 2011) was used to perform an automated meta-analysis of 24 fMRI studies of navigation, revealing common activation across these studies in precuneus, retrosplenial cortex, parahippocampal cortex, and hippocampus. Map is thresholded with a false discovery rate of 0.05.




Figure 2.2 Map and example stimuli used in the experiment. A) Map of the 8 locations (i.e., intersections) on the University of Pennsylvania campus. B) For each intersection, 17 photographs were taken facing each of the cardinal directions (i.e., North, East, South, West).

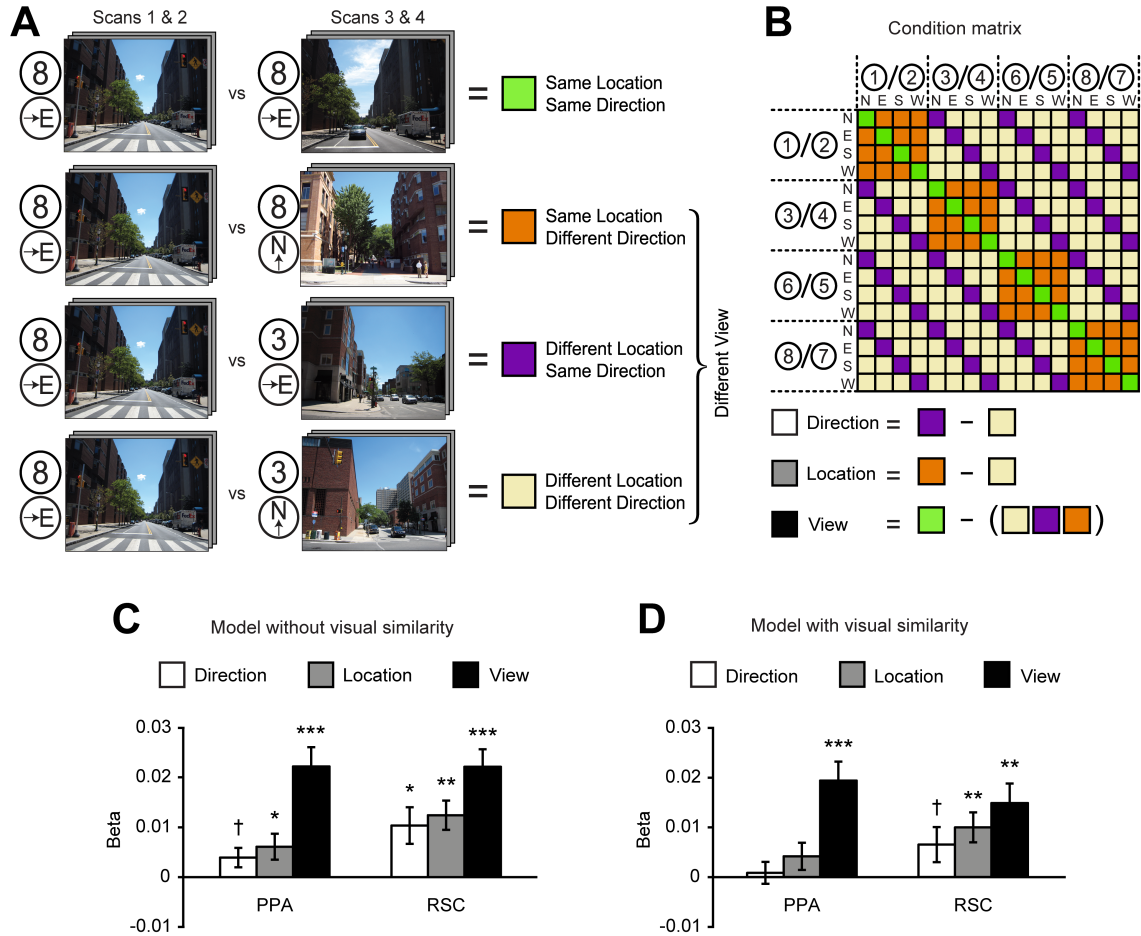


Figure 2.3 Coding of spatial quantities in multivoxel patterns in PPA and RSC. **A)** Multivoxel pattern correlations between each pair of views were sorted into 4 groups based on shared spatial information between the views. Here we show example pairings for each group. **B)** Condition matrix showing all pairings between views, with assignment to groups. Numbered circles refer to campus locations in Figure 2; locations to the left of the slash were shown in the first half of the experiment and locations to the right of the slash were shown in the second half of the experiment. Categorical regressors for direction, location, and view were created by contrasting these conditions as shown. Each regressor was used in a separate multiple regression analysis (see Methods). **C)** When conditions were compared directly without controlling for visual similarity, multivoxel patterns in PPA distinguished between locations and also distinguished between views, and multivoxel patterns in RSC distinguished between directions, locations, and views. Bars represent the difference in pattern similarity between the conditions that constitute each categorical regressor, as measured by the beta weight on that regressor in the multiple regression; error bars are \pm SEM. **D)** When visual similarity was controlled using a parametric regressor, multivoxel patterns in PPA distinguished between views, but not locations or directions, whereas multivoxel patterns in RSC distinguished between both views and locations. *** $P < 0.001$; ** $P < 0.01$; * $P < 0.05$; † $P < 0.10$.

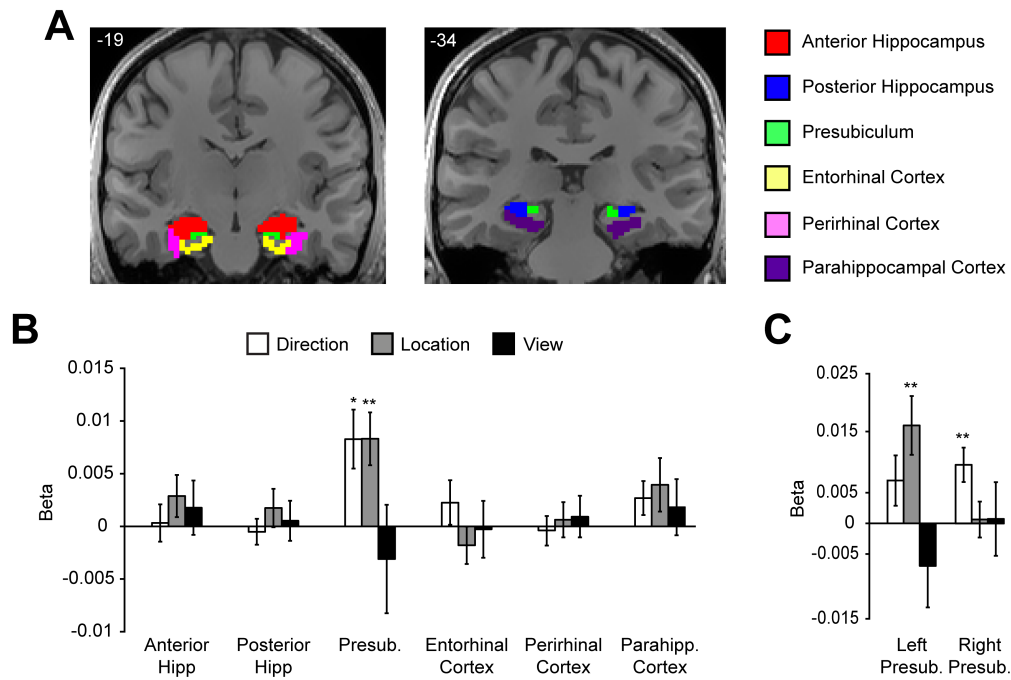


Figure 2.4 Coding of spatial quantities in multivoxel patterns in the medial temporal lobe. **A)** Six anatomical regions were defined in the medial temporal lobes of each subject as described in Methods. The six regions from one subject are displayed on two coronal slices. **B)** Multivoxel patterns in presubiculum distinguished between directions and between locations. No other region showed coding of spatial quantities. **C)** Data for presubiculum, shown separately for each hemisphere, which suggests a difference in spatial coding across hemispheres. Left presubiculum distinguished between locations whereas right presubiculum distinguished between directions. ** $P < 0.01$; * $P < 0.05$.

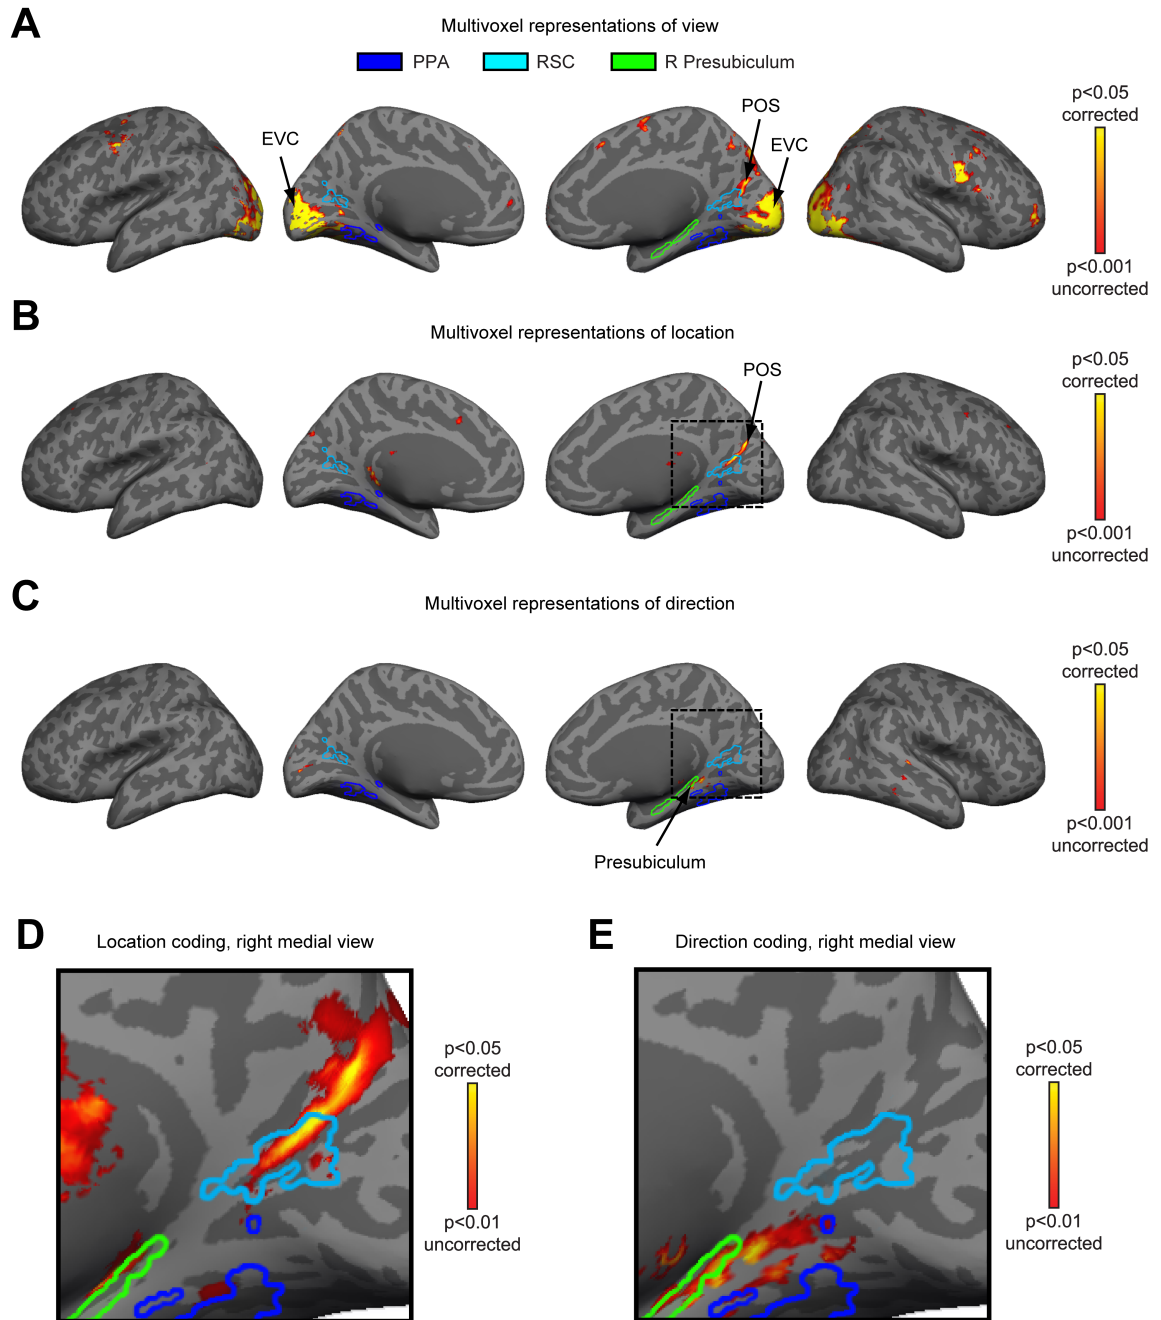


Figure 2.5 Whole-brain searchlight analyses of multivoxel coding of view, location, and direction information. **A)** Views of campus were distinguishable in early visual cortex (EVC) and parietal-occipital sulcus (POS). Results are plotted on the inflated surface of one participant's brain, where dark gray represents sulci and light gray represents gyri. Voxels in yellow are significant ($p<0.05$) after correcting for multiple comparisons via Monte Carlo simulation. Outlines for PPA and RSC were created by calculating the across-subject ROI intersection that most closely matched the average size of each ROI. Outline of right presubiculum was created based on the anatomy of

the participant's brain used for visualization. **B)** Locations were distinguishable in RSC and POS. **C)** Directions were distinguishable in right presubiculum. **D)** Enlarged view of location coding within the portion of the right medial surface indicated by the dotted lines in B. Data are presented at a lower threshold than in **B** to show the extent of location coding. Note that location coding is evident in RSC at this threshold. **E)** Enlarged view of direction coding within the portion of the right medial surface indicated by the dotted lines in C, showing that direction information was primarily present in the anterior calcarine sulcus, and posterior presubiculum. Same threshold as in **D**.




Figure 2.6 Direction tuning curves for right presubiculum. Beta weights, reflecting the average correlation between multivoxel patterns elicited by views in one scan run and patterns elicited by views in the complementary scan run, are plotted as a function of the facing directions of the views. For example, for a given North view, the average correlation was calculated between the pattern for that view and the pattern for all other views facing to the North, East, South, and West (excluding pairs of views obtained at the same intersection); these values were then grouped by direction to give average correlation values for North-North, North-East, North-South, and North-West. Note that there are peaks at North-North, East-East, and South-South, indicating that right presubiculum distinguishes same vs. different direction for North, East and South.

A Thirty-two views of campus



B Seventeen exemplars of ① facing West



Figure 2.7 A) Photographs of each of the 32 views used in the experiment. Note that there were 17 total photographs of each view (1 photograph of each shown). B) Seventeen exemplar photographs for one view, Location 1 facing West.

CHAPTER 3: NEURAL REPRESENTATIONS UNDERLYING REAL-WORLD SPATIAL MEMORY RETRIEVAL

3.1 Abstract

Humans can be cued to retrieve spatial memories by perceiving a visual scene during active navigation or by imagining the relationships between non-visible landmarks, such as when providing directions. However, it is not known whether accessing spatial memory in different ways elicits the same spatial representations in the brain. To test this, we scanned university students with fMRI while they performed two versions of a judgment of relative direction (JRD) task that required them to imagine themselves facing a particular direction at a particular campus location and then report whether a target landmark would be to their left or right. In one version, subjects oriented by imagining the directional relationship between two landmarks, indicated by word cues ("you are at X facing Y"). In the other version, subjects oriented based on a photograph depicting the view from X facing Y. We examined similarities between multivoxel patterns across the two versions of the task to test for representations of the starting location and heading that were independent of the cues to access spatial memory. Entorhinal cortex (ERC) and the scene-selective retrosplenial complex (RSC) coded for heading in a way that abstracted across locations and tasks, but whereas ERC coded for all headings, RSC only coded for the reference direction of the environment in spatial memory. Scene-selective parahippocampal place area (PPA) coded for the starting location in a way that abstracted across headings and tasks. These data suggest that consistent spatial representations are elicited irrespective of whether the spatial memory system is driven by top-down imagery processes or bottom-up perceptual processes, and that at least two distinct heading representations are elicited during JRDs.

3.2 Introduction

The ability to remember the spatial relationships between different places in the world is critical for successful navigation. A navigator must be able to represent allocentric relationships, in which distances and directions are defined relative to an environment-centered reference frame (e.g., “the bookstore is 650 meters west of the ice rink”), as well as egocentric relationships, in which distances and directions are defined relative to the body (e.g., “the bookstore is 300 meters to my right”). These relationships are often retrieved from memory during active navigation, when they may be used to plan a route from the current location to a goal. However, humans can also retrieve the spatial relationships between unseen landmarks at will, such as when describing the route from A to B when currently located at C. It is presumed that both scenarios access the same underlying information even though different processes are invoked to retrieve that information (Byrne et al., 2007). When actively navigating, spatial relationships are retrieved based on visual information available in the current scene and the navigator's recent history of self-motion cues. When providing directions, the current visual scene must be ignored, and spatial relationships must instead be retrieved based on internally generated imagery processes. Do both of these scenarios ultimately access the same spatial representation?

To answer this question, we first consider which brain regions might support these spatial representations. Despite decades of electrophysiological research on rodents, little is known about the neural substrates of complex navigational behaviors such as trajectory calculation and route planning. Prior work has provided evidence for specialized cell types in the medial temporal lobe (MTL) that code for the animal's current location (“place cells” O'Keefe and Dostrovsky, 1971; “grid cells” Hafting et al., 2005) and current heading (“head direction cells” Taube et al., 1990b) during navigation, but it is unknown whether these cells are involved in the navigational behaviors described above. There is no direct evidence that these cells calculate trajectories to distant locations (Navratilova and McNaughton, 2014) and for practical

reasons, it is not known whether these cells represent locations or headings during imagery (although preplay events in place cell populations during navigational decisions suggest that they might; Johnson and Redish, 2007).

To gain traction on these questions, we turn to the human neuroimaging and neuropsychology literatures, which have made considerably more progress in this area and strongly suggest the involvement of parahippocampal and retrosplenial cortex in complex navigational behaviors. The parahippocampal place area (PPA) and retrosplenial complex (RSC) are two regions that are functionally-defined based on their preference for navigational stimuli (i.e., scenes) and are strongly activated during both active (Maguire et al., 1998; Spiers and Maguire, 2006; Op de Beeck et al., 2013) and imagined (Ghaem et al., 1997; Maguire et al., 1997; Rosenbaum et al., 2004; Lambrey et al., 2012) navigation. Whereas PPA is thought to represent the spatial and non-spatial characteristics of the perceived or imagined scene (see Epstein and Vass, 2014 for a review), RSC has been strongly implicated in coding of spatial relationships. In novel environments, the level of activity in RSC predicts the amount of survey knowledge subjects have acquired (Wolbers and Buchel, 2005), and in familiar environments, RSC is strongly engaged during the retrieval of spatial relationships, whether cued by a visual scene (Aguirre and D'Esposito, 1997; Epstein et al., 2007) or by verbal stimuli (Rosenbaum et al., 2004). Although RSC is strongly activated when encoding and retrieving spatial information, this alone does not indicate that it is necessary for representing spatial relationships. The strongest evidence for this claim comes from neuropsychological studies, which have shown that damage to this region leads to a deficit referred to as "heading disorientation," which renders patients unable to retrieve the directional relationships between landmarks despite intact memory for the landmarks' identities (Aguirre and D'Esposito, 1999). The fact that patients cannot access this information when actively navigating or when attempting to imagine routes (Takahashi et al., 1997) suggests that this region is involved in both processes.

Taken together, the neuroimaging and neuropsychology data suggest that RSC plays an important role in the representation of spatial relationships. Recent neuroimaging studies have begun to interrogate the precise representations supported by RSC and suggest that this region can distinguish between allocentric headings (Baumann and Mattingley, 2010; Vass and Epstein, 2013) and between distinct locations in an environment (Vass and Epstein, 2013). However, both of these studies investigated representations only during perception of visual scenes. It is entirely unknown whether RSC represents this information during spatial imagery and if so, whether the representation takes the same form as during perception.

We address these questions in the current experiment by measuring neural activity while subjects performed judgments of relative direction (JRDs). This task requires subjects to retrieve spatial relationships and is commonly used in the behavioral literature to study the reference frames underlying spatial memory (see McNamara, 2003 for a review). On each trial, university students indicated whether a target landmark would be on their left or right given a particular view of campus. Critically, we varied the manner in which this view was conveyed to subjects. In one half of the scan runs, subjects imagined the view based on the relationship between two landmarks (e.g., “at X facing Y”); in the other half of the scan runs, subjects self-localized based on a photograph of campus. We then tested for representations of heading and location that abstracted across perception and imagery by measuring the similarity of multivoxel patterns elicited during the JRDs. To anticipate, we observe abstract representations of heading in RSC and entorhinal cortex and abstract representations of location in PPA.

3.3 Materials and Methods

SUBJECTS. Sixteen healthy right-handed subjects (9 female, mean age = 22 ± 0.2 y) with normal or corrected-to-normal vision were scanned with fMRI. All subjects had at least two years of experience with the University of Pennsylvania campus and were

either current undergraduate students or recent graduates of the University of Pennsylvania. Subjects provided written informed consent in compliance with procedures approved by the University of Pennsylvania Institutional Review Board.

PRE-SCREENING SESSION. At least one day before the fMRI scan, subjects were brought in for an extensive pre-screening appointment to confirm their knowledge of the Penn campus. In order to be eligible for the scan session, subjects were required to successfully complete five computerized tests of spatial knowledge.

First, we assessed subjects' familiarity with the 77 campus landmarks used in the experiment. For each landmark, subjects were presented with the name of the landmark and asked to 1) indicate whether they were familiar with that place (yes/no), 2) rate how vividly they could imagine that place using a 5 point scale, and 3) rate their confidence that they could navigate to that place from their home using a 4 point scale.

Second, we tested subjects' explicit knowledge of the allocentric directional relationships between campus landmarks. On each trial, subjects were presented with a statement of the form "X is _____ of Y," where X and Y were the names of landmarks, and subjects indicated whether the correct directional relationship was North (N), East (E), South (S), or West (W). This test consisted of 72 trials total, 16 trials which queried the converse of directional relationships that subjects would later retrieve in the fMRI experiment (e.g., subjects responded "X is North of Y" and were later asked in the scanner to imagine standing at X facing Y, i.e., facing South), and 56 trials whose directional relationships were not probed in the scanner.

Third, subjects trained on the two versions of the judgment of relative direction (JRD) task that were used in the fMRI experiment. In the verbal version of the JRD task, subjects were presented with the names of three campus landmarks (X, Y, Z) and were asked to imagine they were at X facing Y and then indicate whether Z would be on their left or right given that imagined viewpoint. In the picture version of the JRD task, subjects were shown a photograph of what it looks like to stand at X facing Y, with the

name of the target landmark (Z) superimposed onto the photograph; they then indicated whether Z would be on their left or right given the view shown in the photograph. Subjects completed 200 practice trials of each version of the task, with no time limit imposed for responding. All practice trials started from one of four locations (statue, bridge, two intersections) and asked subjects to imagine or perceive headings to the N, E, S, or W. These trials were independent from those used in the main fMRI experiment, which consisted of N, E, S, and W views from four distinct campus locations. This procedure allowed subjects to gain practice with the task itself without exposing them to the stimuli of interest.

Fourth, subjects completed 72 trials of each version of the JRD task at the speed of presentation in the fMRI experiment (5.5 seconds + 0.5 second interstimulus interval; ISI); these trials were drawn from the same stimulus set as the unspeeded practice trials.

Finally, to confirm subjects' familiarity with the 77 campus landmarks, they were presented with the name of a campus landmark and were asked to select the matching photograph of that landmark from a deck of cards. The landmarks were grouped into 4 decks: open spaces (e.g., intersections, N=14), academic buildings (N=18), dormitories and recreational facilities (N=26), and restaurants (N=19). Subjects completed all trials corresponding to the first deck before proceeding to the second, third, and fourth decks in turn. Subjects that successfully completed all five tests of spatial knowledge (16/25 subjects) were subsequently scanned with fMRI.

MRI ACQUISITION. Scanning was performed at the Hospital of the University of Pennsylvania using a 3T Siemens Trio scanner equipped with a 32-channel head coil. High-resolution T1-weighted anatomical images were acquired using a three-dimensional magnetization-prepared rapid-acquisition gradient echo pulse sequence [repetition time (TR) = 1620 ms; echo time (TE) = 3 ms; inversion time = 950 ms; voxel size = 0.9766 x 0.9766 x 1 mm; matrix size = 192 x 256 x 160]. T2*-weighted images sensitive to blood oxygenation level-dependent contrasts were acquired using a

gradient-echo echo-planar pulse sequence (TR = 3000 ms; TE = 30 ms; voxel size = 3 x 3 x 3 mm; matrix size = 64 x 64 x 44).

DESIGN. Each subject completed nine scan runs in the following order: three runs of the verbal version of the JRD task, two functional localizer runs to identify scene-selective regions, three runs of the picture version of the JRD task, and a high-resolution anatomical scan. The verbal version of the JRD task was always performed before the picture version to ensure that subjects could not imagine the specific photographic stimuli during the verbal version of the task.

Each of the experimental runs were 10.8 minutes in length and consisted of 93 6-s stimulus trials interspersed with 6 12-s null trials. There was an additional 18-s null trial at the end of each run to capture the hemodynamic response of the final stimulus trial.

On each trial of the verbal version of the JRD task (5.5 s + 0.5 s ISI), subjects viewed a multicolored texture (1024 x 768 pixels) overlaid with the names of three campus landmarks, presented centrally on separate lines and flanked by nonsense characters that extended to the full width of the screen. Subjects were asked to imagine a particular view of campus indicated by the first two landmarks (i.e., at X facing Y) and indicate via button press whether the third landmark would be on their left or right given that imagined view.

On each trial of the picture version of the JRD task (5.5 s + 0.5 s ISI), subjects viewed a photograph of a view on the Penn campus (1024 x 768 pixels) overlaid with the name of a centrally presented target landmark. Subjects indicated via button press whether the target landmark would be on their left or right given the depicted view of campus.

In both versions of the task, subjects made a “left” response by pressing a button with their left thumb and a “right” response by pressing a button with their right thumb. Responses from the verbal runs of one participant were not collected due to a

technical error. Null trials consisted of 12 s of a gray screen, during which subjects made no response.

Trials were ordered using a continuous carryover sequence (Aguirre, 2007) which fully counterbalanced the trials at the level of the JRD's campus view (e.g., Location 1 facing N) such that each view was presented before and after every other view including itself exactly once. Subjects completed one full carryover sequence for each task, which was spread across three scan runs. Because this created an interruption of the carryover sequence between runs, each run began with the last 3 trials from the end of the previous run (or in the case of the first run, the last 3 trials from the end of the third run), which were subsequently removed and not analyzed.

Over the course of the experiment, subjects were asked to imagine or view 16 different views of campus, corresponding to views facing the four cardinal directions (N, E, S, W) at four campus locations: two statues, a courtyard, and a large compass inlaid in a walkway. Each of these views was paired with 2 unique target landmarks, one to the left and one to the right, for a total of 32 JRD problems. Because each view was presented 17 times over the course of 3 runs, subjects solved each JRD problem 8 or 9 times. The ordering of the 2 JRD problems within each view condition was randomized with the constraint that subjects never solved the same problem on adjacent trials. Target landmarks were selected so that their relative bearing was within 70 degrees of 90 or 270 degrees (mean deviation from 90 or 270 = 19 ± 3 degrees). The same set of target landmarks was used in both versions of the task. In the verbal version of the task, each location was paired with 4 unique landmarks that served as the N, E, S, and W direction cues. Direction cues were selected so that the bearing was as close as possible to the cardinal directions as defined by the campus grid (mean deviation = 10 ± 1 degrees). All target landmarks and direction cues were unique with one exception: the S-facing cue at Location 3 was also a target landmark when facing E at Location 3.

In addition to the experimental runs, subjects also completed two functional localizer runs (5.25 min each), during which subjects performed a one-back repetition task. Stimuli were presented for 800 ms with 200 ms ISI and consisted of images of scenes, objects, and scrambled objects, which were presented in 15-s blocks.

STIMULI. Stimuli were designed to minimize visual similarities across the two versions of the task. Verbal stimuli consisted of three lines of text consisting of the names of three campus landmarks, presented overlaid on a colorful texture (1024 x 768 pixels). Landmark names were presented in one of 17 fonts and were flanked by nonsense characters that extended to the full width of the image. The background images consisted of 17 phase-scrambled versions of each of 17 different textures (289 images total), which were never repeated during the experiment. Picture stimuli consisted of a photograph of a view of campus overlaid with a centrally presented landmark name. Note that because the landmark name appears in both tasks, there is cross-task visual similarity for the SL-SH-SR condition. Landmark names were presented in 17 different fonts from those used in the verbal version. There were 17 distinct photographs of each of the 16 views, which were never repeated during the experiment. We randomized the assignment of fonts and images separately for each subject such that each view was paired with each font and each image (17 photographs of the view or 17 different textures) exactly once.

DATA PREPROCESSING. Functional images were corrected for differences in slice acquisition timing using VoxBo's sliceacq function (<http://www.nitrc.org/projects/voxbo>), which resamples slices in time using sinc interpolation to match the first slice of each volume. Data from each scan were then preprocessed using FSL's FEAT 5.98 (Jenkinson et al., 2012), which included prewhitening to account for autocorrelation in time, high pass temporal filtering at a period of 100 s, and motion correction using MCFLIRT (Jenkinson et al., 2002). Outlier volumes were identified using the Artifact

Detection Tools (http://www.nitrc.org/projects/artifact_detect) and defined as volumes with global signal values more than 3.5 SD away from the mean or volumes in which subject motion exceeded 3 mm. Data from the functional localizer scans were smoothed with a 5-mm full-width-half-maximum (FWHM) Gaussian filter; data from the experimental runs were not smoothed. We discarded the first three trials (i.e., six volumes) of each experimental scan run as these trials served to re-instantiate the continuous carryover sequence from the previous scan run.

FUNCTIONAL REGIONS OF INTEREST. Data from the functional localizer scans were used to identify two scene-responsive regions, the parahippocampal place area (PPA) and retrosplenial complex (RSC), which have been previously implicated in place recognition and navigation. PPA and RSC were defined for each subject using a contrast of scenes > objects and a group-based anatomical constraint of scene-selective activation derived from 42 independent subjects previously scanned by our lab (Julian et al., 2012). For each hemisphere of PPA and RSC, we selected the 100 voxels within the group-based mask that showed the strongest scenes > objects effect. This method defines regions in a threshold-free manner and ensures that ROIs can be defined in both hemispheres for all subjects.

ANATOMICAL REGIONS OF INTEREST. We anatomically defined six regions within the medial temporal lobe: anterior hippocampus, posterior hippocampus, presubiculum, entorhinal cortex, perirhinal cortex, and parahippocampal cortex. The hippocampus and presubiculum were defined using the fully automated segmentation protocol in FreeSurfer 5.1 (Van Leemput et al., 2009). This technique uses Bayesian inference on an upsampled version of the T1 structural image to determine the likely hippocampal subfield identity of each 0.5 x 0.5 x 0.5 mm voxel. We first assigned a subfield identity to each “mini-voxel” by selecting the subfield with the highest probability. We then assigned the identity of each 3 x 3 x 3 mm functional voxel according to the most

commonly occurring subfield across the 216 mini-voxels. The hippocampus ROI was defined as the union of the CA1, CA2/3, CA4/Dentate Gyrus, and subiculum subregions. We then divided the hippocampus into anterior and posterior subregions at the middle coronal slice of each subject's hippocampus. Entorhinal, perirhinal, and parahippocampal cortices were defined based on manual parcellation of the T1 anatomical image in ITK-SNAP (<http://www.itksnap.org>) following the protocol in Pruessner et al. (2002).

GENERAL LINEAR MODELS. All models were run using FEAT 5.98 in each subject's native space. For the experimental runs, there were six first-level models, one for each scan run, and two higher-level models that computed the fixed effects of the three scan runs from each task. Each experimental model contained 32 regressors of interest, one for each JRD problem, which modeled the presentations of a given JRD as a 6-s boxcar convolved with a double gamma hemodynamic response function and were high pass filtered at a period of 100 s. First-level models also contained nuisance regressors corresponding to the six motion parameters calculated by MCFLIRT and stick functions for any volumes identified as outliers by the Artifact Detection Tools.

For the localizer models, there were two first-level models, one for each scan run, and one higher-level model that computed the fixed effects of the two scan runs. Each model contained three regressors of interest, one for each condition (scenes, objects, scrambled objects), which modeled the blocks of a given condition as a 15-s boxcar convolved with a double gamma hemodynamic response function and were high pass filtered at a period of 100 s. Models contained the same nuisance regressors as described for the experimental runs.

CROSS-TASK DECODING OF LOCATION AND HEADING. To determine whether each ROI contained abstract information about location or heading, we calculated the cross-task correlations between the multivoxel patterns elicited in the two tasks (Haxby et al.,

2001). First, for each ROI, we extracted the voxelwise t statistics for the 32 JRDs calculated in each task's higher-level general linear model (i.e., 64 patterns total for each ROI). Second, we separately normalized the patterns from each task by calculating the mean pattern across JRDs from that task and subtracting this mean pattern from each of the 32 individual JRD patterns. Third, we created a 32 x 32 correlation matrix by calculating the cross-task Pearson correlation between all pairs of JRDs. Each cell of the correlation matrix belonged to 1 of 8 possible groups based on whether the pair of JRDs shared the same location (Same Location, SL; Different Location, DL), heading (Same Heading, SH; Different Heading, DH) or response (Same Response, SR; Different Response, DR): 1) SL-SH-SR (i.e., same JRD across tasks); 2) SL-SH-DR (e.g., Location 1 Heading N Left Response vs. Location 1 Heading N Right Response); 3) SL-DH-SR (e.g., Location 1 Heading N Left Response vs. Location 1 Heading E Left Response); 4) SL-DH-DR (e.g., Location 1 Heading N Left Response vs. Location 1 Heading E Right Response); 5) DL-SH-SR (e.g., Location 1 Heading N Left Response vs. Location 2 Heading N Left Response); 6) DL-SH-DR (e.g., Location 1 Heading N Left Response vs. Location 2 Heading N Right Response); 7) DL-DH-SR (e.g., Location 1 Heading N Left Response vs. Location 2 Heading E Left Response); 8) DL-DH-DR (e.g., Location 1 Heading N Left Response vs. Location 2 Heading E Right Response). Finally, we submitted these 8 groups of correlations to a Location (Same, Different) x Heading (Same, Different) x Response (Same, Different) repeated measures ANOVA to test whether pattern similarity varied based on any of these three factors.

SEARCHLIGHT ANALYSES FOR CROSS-TASK DECODING OF LOCATION AND HEADING. To test for cross-task coding of location or heading across the entire brain, we implemented a searchlight procedure (Kriegeskorte et al., 2006) to perform pattern analyses in small spherical ROIs (radius = 5 mm) centered on every voxel of the brain in turn. The procedure was identical to that described above up until the point of the ANOVA. Instead of the ANOVA, we calculated the main effect of location restricted to

different headings ($\text{mean}(\text{SL-DH-SR}, \text{SL-DH-DR}) - \text{mean}(\text{DL-DH-SR}, \text{DL-DH-DR})$) and the main effect of heading restricted to different locations ($\text{mean}(\text{DL-SH-SR}, \text{DL-SH-DR}) - \text{mean}(\text{DL-DH-SR}, \text{DL-DH-DR})$). This allowed us to identify regions that code for location in a way that abstracts across tasks and headings and regions that code for heading in a way that abstracts across tasks and locations. These mean correlation differences were assigned to the center voxel of the spherical ROI, generating 2 whole-brain maps for each subject. Individual subject maps for each contrast were then transformed into standard space using FLIRT 5.5 (Jenkinson and Smith, 2001; Jenkinson et al., 2002) and smoothed with a 9-mm FWHM Gaussian filter before performing separate higher-level random effects analyses to identify voxels that reliably coded for location or heading across subjects. Finally, to estimate the true Type 1 error rate for each type of spatial coding, we performed permutation testing (Nichols and Holmes, 2002) using FSL's randomise function with 12-mm variance smoothing and 10,000 permutations per contrast. We report voxels that are significant at $P < 0.05$, corrected for multiple comparisons across the entire brain.

REPRESENTATIONAL STRUCTURE ANALYSES. To characterize the structures of the ROI representations, we performed three analyses. First, to visualize the representational similarity of the 32 JRDs, we submitted a modified version of the 32 x 32 cross-task correlation matrix to multidimensional scaling (MDS). We converted the correlations to correlation distances by subtracting each correlation from 1 (i.e., $1 - r$), set the diagonal of the matrix to zero (since MDS does not take into account the pattern stability of the individual JRDs), and then averaged across the upper and lower triangles of the matrix to make it symmetric. We calculated the mean distance matrix by averaging together the matrices of the 16 subjects and visualized the relationships in this distance matrix using the MDSConditions function from the Representational Similarity Analysis (RSA) Toolbox (Nili et al., 2014). This procedure produces a plot of a 2-dimensional plane in which each JRD is represented by a point, and distances

between JRDs on the plot correspond to distances between multivoxel patterns. We made two versions of each plot, which were identical except for the coloring of the points, and were colored either according to the starting location of the JRD or the heading of the JRD. This allowed for visual inspection of any clustering of points based on location or heading.

Second, we created a 4 x 4 location correlation matrix and a 4 x 4 direction correlation matrix to quantify the relationships observed with MDS. To generate these matrices, we took the original 32 x 32 correlation matrices for each subject, set the diagonal to zero, and averaged the matrix across subjects. Thus, this matrix differed from the one submitted to MDS in that it contained correlations rather than distances and was asymmetric, which afforded us the ability to visualize asymmetric representational relationships (e.g., verbal N trials could be similar to picture E trials even if picture N trials were not similar to verbal E trials). We then calculated the mean correlation for each location pair (N=16) and each direction pair (N=16) across tasks and inserted them into the respective 4 x 4 matrices. Elements on the diagonal of the matrix represent cross-task correlations for the same location (e.g., Location 1 - Location 1) or the same direction (e.g., N-N), and elements on the off-diagonal of the matrix represent cross-task correlations for different locations (e.g., Verbal Location 1 - Picture Location 2) or different directions (e.g., Verbal N - Picture E).

Third, we compared the magnitude of coding across the four directions or four locations using repeated measures ANOVAs. We calculated 8 values from the correlation matrices: 4 within-direction (e.g., N-N) or within-location correlations corresponding to the diagonal of the matrix, and 4 between-direction (e.g., mean[N-E, N-S, N-W, E-N, S-N, W-N]) or between-location correlations corresponding to off diagonal rows and columns. These values were submitted to Direction (N, E, S, W) x Similarity (Within-Direction, Between-Direction) and Location (1, 2, 3, 4) x Similarity (Within-Location, Between-Location) repeated measures ANOVAs. All reported t tests are two-tailed.

3.4 Results

BEHAVIORAL PERFORMANCE. During the fMRI experiment, subjects performed two versions of a judgment of relative direction (JRD) task, in which they reported via button press whether a target landmark would be on their left or right given a particular view of campus (Figure 3.1). The critical difference between the two versions of the task is the manner in which the view of campus is conveyed to the subject. In the verbal version of the task, subjects were presented with three lines of text, each of which contained the name of a landmark to be used as the starting location, heading cue, or target respectively. Thus, in order for subjects to determine their allocentric heading in the environment, they had to solve the directional relationship between the first two landmarks (e.g., “Meyerson Hall is East of Broken Button, so facing East”). In the picture version of the task, subjects were presented with the name of the target landmark overlaid on a photograph of a view of campus. Subjects had to determine both the starting location and heading using the visual information available in the photograph. This was likely achieved by either relating the visible scene to similar views stored in memory (e.g., “this is what it looks like facing East at Broken Button”) or inspection of the spatial relationships between features visible in the scene (e.g., “Woodland Walk is headed diagonally to the left and I can see a red statue in the distance, so I must be facing East at Broken Button”). Subjects performed both versions of the task rapidly and accurately, which was expected given the extensive pre-screening and the pre-scan practice session which used an independent set of stimuli (see Methods for details). Subjects were equally accurate on both versions of the task (% Correct Verbal: 93.6 ± 0.02 ; % Correct Picture: 96.2 ± 0.01 ; $t(14) = 1.4$, $P = 0.19$), but responded significantly faster during the Picture version of the task (Verbal: 2.9 ± 0.1 s; Picture: 1.8 ± 0.1 s; $t(14) = 8.3$, $P < 0.001$), even though this version of the task required them to solve an additional spatial variable (i.e., location).

We next tested whether reaction times varied as a function of either starting location or allocentric heading, with faster reaction times indicating that the representation is more easily accessible in spatial memory (Shelton and McNamara, 1997; Montello et al., 2004). In both versions of the task, reaction time was significantly modulated by starting location (one-way repeated-measures ANOVA; Verbal: $F(3,42)=3.2$, $P=0.03$; Picture: $F(3,45)=3.6$, $P=0.02$) and by allocentric heading (Verbal: $F(3,42)=6.9$, $P=0.001$; Picture: $F(3,45)=21.9$, $P<0.00001$). Post-hoc *t*-tests revealed that in both versions of the task, subjects were generally faster for judgments that started at Location 1, although not all pairwise comparisons were significant (Verbal: Location 1 vs. 2 $t(14)=2.2$, $P=0.046$; Verbal Location 1 vs. 3 $t(14)=1.9$, $P=0.08$; Verbal Location 1 vs. 4 $t(14)=3.6$, $P=0.003$; Picture Location 1 vs. 2 $t(15)=2.6$, $P=0.02$; Picture Location 1 vs. 3 $t(15)=3.2$, $P=0.007$; Picture Location 1 vs. 4 $t(15)=0.7$, $P=0.47$). Subjects were also significantly faster for JRDs headed North (N) than JRDs headed East (E), South (S), or West (W; all $P_s<0.02$); in the Picture version, subjects were also faster for S than E ($t(15)=2.3$, $P=0.03$). Thus, although all of our subjects had long-term, real world experience with the environment, their behavioral performance indicated a privileged representation of N-facing headings and, to a lesser extent, a particular location on campus.

Our last analysis of the reaction time data was a test of whether subjects exhibited behavioral priming for either the starting location or allocentric heading across successive JRD trials (Figure 3.2). We tested whether the reaction time to a particular trial was faster when preceded by a trial that shared either the same starting location or same allocentric heading. We observed significant location priming in the Verbal version of the task ($t(14)=6.3$, $P=0.00002$; individual locations all $P_s < 0.04$), but not the Picture version ($t(15)=0.2$, NS; individual locations all $P_s > 0.4$). However, location priming in the Verbal version should be interpreted with caution since it may simply reflect faster processing of the visual stimulus due to repetition of the name of the starting location across trials. There was not an overall effect of direction priming in

either version of the task (Verbal: $t(14)=0.6$, $P=0.59$; Picture: $t(15)=0.6$, $P=0.58$), but there was priming for N in the Verbal version ($t(14)=3.2$, $P=0.007$).

CROSS-TASK DECODING OF LOCATION AND HEADING. The primary goal of this experiment was to identify brain regions whose multivoxel activity patterns contained information about particular real world locations or allocentric headings in a way that abstracts across the two versions of the task, which cued spatial memory with different types of stimuli (i.e., verbal or photograph) and required different cognitive processes to access the memory (i.e., spatial imagery or visual inspection of the scene). We hypothesized that if a brain region contains information about a particular spatial quantity (e.g., allocentric heading), then 2 JRD problems that share that spatial quantity (e.g., both heading N) should elicit activity patterns that are more similar than 2 JRD problems that differ on that spatial quantity (e.g., heading N vs. heading E). To test this, we measured the pattern similarity (i.e., Pearson correlation) between all pairs of JRDs across tasks (Figure 3.3). We then grouped these pattern similarities based on three factors: whether the pair of JRDs were from the *Same* or *Different Location* (SL or DL), faced the *Same* or *Different Heading* (SH or DH), and elicited the *Same* or *Different Response* (i.e., target in the same egocentric direction; SR or DR). We then submitted these 8 groups of correlations to a Location (Same, Different) x Heading (Same, Different) x Response (Same, Different) repeated-measures ANOVA to investigate whether pattern similarity within a region of interest (ROI) was modulated by any of these factors.

We first report results from two functionally-defined scene-selective regions, Retrosplenial Complex (RSC) and Parahippocampal Place Area (PPA; Figure 3.3). Based on previous work (Vass and Epstein, 2013), we predicted that RSC would be sensitive to the spatial quantities associated with the JRD problems whereas PPA would be sensitive to features of the local scene, whether that scene was perceived or imagined. Pattern similarities in RSC were significantly greater for JRDs that shared the

same allocentric heading (main effect of heading: $F(1,15)=12.7$, $P=0.003$), an effect that was modulated by response (interaction between heading and response: $F(1,15)=7.2$, $P=0.02$) such that pattern similarity was higher for JRDs that elicited the same response, but only when they also shared the same heading. Pattern similarity in RSC was not significantly modulated by location (main effect: $F(1,15)=2.5$, $P=0.14$). Pattern similarities in PPA were significantly greater for JRDs that shared the same location (main effect: $F(1,15)=9.0$, $P=0.009$) or the same heading (main effect: $F(1,15)=4.5$, $P=0.05$), and these effects were modulated by response (interaction between location and response: $F(1,15)=14.5$, $P=0.002$; interaction between location, heading, and response: $F(1,15)=11.8$, $P=0.004$).

At first glance, the finding that PPA pattern similarity was modulated by location and heading would appear to be counter to our hypothesis. However, a cursory examination of the mean pattern similarities by correlation type (Figure 3.3) shows that PPA exhibited very high correlations for the same JRD across tasks (i.e., SL-SH-SR correlation type). Indeed, pattern similarity was significantly higher for this correlation type than for any of the other seven correlation types (2-tailed t -tests; all $P_s < 0.01$). To test whether the SL-SH-SR condition could have driven either the main effect of location or heading in PPA, we excluded this condition and performed separate 2x2 repeated-measure ANOVAs on Location x Response, restricted to different headings only, and Heading x Response, restricted to different locations only. Under these conditions, PPA still exhibited a main effect of location ($F(1,15)=12.4$, $P=0.003$), but the main effect of heading was no longer significant ($F(1,15)=2.3$, $P=0.15$); there were no effects of response (all $P_s > 0.24$). When we ran the same Heading x Response ANOVA in RSC, there was still a significant main effect of Heading ($F(1,15)=6.9$, $P=0.02$), but there was no longer an interaction with response ($F(1,15)=0.06$, $P=0.81$). Thus, the results indicate that patterns in RSC coded for allocentric heading across stimulus types whereas patterns in PPA coded for starting location across stimulus types. To test whether these effects significantly differed between regions, we computed a location

index for each region, defined as [Same Location (mean of SL-DH-SR & SL-DH-DR) - Different Location (mean of DL-DH-SR & DL-DH-DR)], and a heading index for each region, defined as [Same Heading (mean of DL-SH-SR & DL-SH-DR) - Different Heading (mean of DL-DH-SR & DL-DH-DR)]. Two-tailed *t*-tests comparing the indices across regions showed that location coding was significantly stronger in PPA than RSC ($t(15)=3.6$, $P=0.003$), but the difference in heading coding between RSC and PPA was not significant ($t(15)=1.7$, $P=0.11$).

We next consider pattern similarities within seven anatomically-defined regions within the medial temporal lobe (MTL): anterior and posterior hippocampus, left and right presubiculum, entorhinal cortex (ERC), perirhinal cortex (PRC), and parahippocampal cortex (PHC). These regions were selected based on prior human (Ekstrom et al., 2003; Hassabis et al., 2009; Doeller et al., 2010; Jacobs et al., 2010; Jacobs et al., 2013; Miller et al., 2013; Vass and Epstein, 2013) and animal (O'Keefe and Dostrovsky, 1971; Taube et al., 1990b; Georges-Francois et al., 1999; Robertson et al., 1999; Cacucci et al., 2004; Hafting et al., 2005; Boccara et al., 2010) studies of spatial memory, which have implicated these regions in coding of spatial quantities.

We submitted pattern similarities from each of these regions to Location x Heading x Response ANOVAs and report effects using a Bonferroni-corrected threshold of $P<0.007$ (Figure 3.4). At this conservative threshold, we observed only one significant main effect: pattern similarities in ERC were greater for JRDs that shared the same heading than JRDs with different headings ($F(1,15)=11.2$, $P=0.004$). ERC also exhibited a trend for the interaction between location and heading ($F(1,15)=6.2$, $P=0.03$) such that heading coding was stronger for JRDs from the same location than JRDs from different locations. We also observed a significant interaction between location and response in left presubiculum ($F(1,15)=10.2$, $P=0.006$): JRDs from the same location showed greater pattern similarity when they elicited the same response, but JRDs from different locations showed greater pattern similarity when they elicited different

responses. There were no main effects in left presubiculum, even at more lenient thresholds (all P s > 0.16). In contrast to our previous work (Vass and Epstein, 2013), we did not observe a main effect of heading in either left or right presubiculum (P s > 0.73). At a less stringent threshold of $P < 0.05$, we observed coding of location in PHC ($F(1,15)=5.2$, $P=0.04$) and right presubiculum ($F(1,15)=5.0$, $P=0.04$), though the presubiculum effect was in the opposite direction, such that JRDs from different starting locations had greater pattern similarity than JRDs from the same starting location. In sum, when we interrogated spatial coding within the MTL, we observed significant coding of heading in ERC and weaker sensitivity to location in PHC and right presubiculum.

Finally, to test whether regions outside of our pre-defined ROIs showed multivoxel coding of starting location or allocentric heading, we performed searchlight analyses (Kriegeskorte et al., 2006), in which we calculated the main effect of location and the main effect of allocentric heading from the restricted 2x2 ANOVAs in spherical regions centered on every voxel of the brain in turn (Figure 3.5). We used the restricted two-way ANOVAs rather than the three-way ANOVA because the SL-SH-SR condition does have shared visual information across tasks (i.e., the target word) and excluding this condition lead to different results in at least one visually-responsive region, PPA. The searchlight analysis for the main effect of heading revealed a cluster in left medial parietal cortex (-9 , -63 , 21) and smaller clusters in left superior frontal gyrus (-15 , 33 , 48) and left middle temporal gyrus (-51 , -24 , -21). No region showed a main effect of location at levels exceeding the permutation-corrected threshold (Nichols and Holmes, 2002).

To summarize, we observed that specific regions in medial temporal and medial parietal cortex coded for spatial quantities in a way that generalized across the nature of the stimulus and the cognitive processes required to access long-term spatial memory. In particular, we observed abstract coding of location in PPA and abstract

coding of heading in RSC and ERC. We now present analyses that aim to characterize the location and heading codes that we observed in these regions.

REPRESENTATIONAL STRUCTURE OF RSC, PPA, AND ERC. To better understand the representations in RSC, PPA, and ERC, we visualized the between-JRD correlations using multidimensional scaling (MDS). For each ROI, we converted each subject's cross-task correlation matrix to a distance matrix ($1 - r$), averaged the distance matrices across subjects, and performed MDS using the Representational Similarity Analysis Toolbox (Nili et al., 2014). The resulting plots provide an intuitive visual representation of the data where each of the 32 JRDs is plotted as a point on a two-dimensional plane and physical distance between points reflects correlation distance between representations. Thus, JRDs with similar representations will occupy nearby locations on the MDS plot. We color coded the points based on either location and heading in order to visualize whether JRDs clustered in the MDS plots based on these factors (Figures 6A, 7A; N.B. unlike the ANOVAs, the MDS analysis does not contain categorical information).

Although the ANOVA results for RSC only provided evidence for coding of heading, the MDS plots appeared to show grouping of JRDs by both heading and location: JRDs headed North were grouped together and JRDs starting at Location 1 were grouped together. JRDs in PPA also appeared to show moderate clustering by heading and location. JRDs in ERC were clustered by heading for East, South, and West, and did not appear to be clustered by location. We next quantified the representational structures we observed in the MDS plots.

In order to examine coding of each direction and location separately, we quantified the average between-JRD similarity (i.e., off diagonal elements of the correlation matrix; this excludes the SL-SH-SR condition) for each direction pair and each location pair and plotted them in matrix form (Figures 3.6B, 3.7B). The diagonal elements of each matrix represent the mean cross-task correlation for the same

direction (e.g., N-N) or location (e.g., Location 1 - Location 1). The off-diagonal elements represent the mean cross-task correlation for different directions (e.g., N-E) or locations (e.g., Location 1 - Location 2). For each direction and location, we calculated the mean between-direction or between-location correlation by averaging together the appropriate off-diagonal elements of the matrix (e.g., N Between = average of N-E, N-S, N-W, E-N, S-N, W-N). The resulting correlations were then submitted to Direction (N, E, S, W) x Similarity (Within-Direction, Between-Direction) and Location (1, 2, 3, 4) x Similarity (Within-Location, Between-Location) repeated measures ANOVAs to test whether the magnitude of coding differed across directions or locations respectively, as indicated by a significant interaction effect (Figures 3.6C, 3.7C). We next describe the ANOVA results for RSC, PPA, and ERC.

Consistent with the MDS plot, RSC exhibited coding of heading which differed by direction (main effect of similarity: $F(1, 15)=10.1$, $P=0.006$; main effect of direction: $F(3, 45)=3.9$, $P=0.01$; interaction effect: $F(3, 45)=3.9$, $P=0.01$) and was stronger for JRDs headed N than JRDs headed E ($t(15)=2.2$, $P=0.047$), S ($t(15)=2.2$, $P=0.04$), or W ($t(15)=3.2$, $P=0.006$; all other P s > 0.68). The Location x Similarity interaction was not significant in RSC ($F(3, 45)=1.7$, $P=0.17$), and the specific t -test for Location 1 was marginal ($t(15)=1.9$, $P=0.08$). PPA exhibited equivalent coding of all locations (main effect of similarity: $F(1, 15)=11.0$, $P=0.005$; main effect of location: $F(3, 45)=0.8$, $P=0.51$; interaction effect: $F(3, 45)=0.9$, $P=0.45$), but did not code for heading (main effect of similarity: $F(1, 15)=1.5$, $P=0.24$; main effect of direction: $F(3, 45)=2.0$, $P=0.13$; interaction effect: $F(3, 45)=2.0$, $P=0.12$). Although the direction matrix for PPA suggested possible coding of N, the specific t -test was only marginal ($t(15)=1.9$, $P=0.07$). Unlike RSC, the magnitude of heading coding did not differ between directions in ERC (main effect of similarity: $F(1, 15)=7.2$, $P=0.02$; main effect of direction: $F(3, 45)=1.0$, $P=0.41$; interaction effect: $F(3, 45)=1.9$, $P=0.15$). Although only the specific t -tests for E and W were significant (N: $t(15)=0.5$, $P=0.59$; E: $t(15)=2.6$, $P=0.02$; S: $t(15)=0.4$, $P=0.66$; W: $t(15)=2.8$, $P=0.01$), there were no significant differences between directions, though E

tended to be stronger than N ($t(15)=1.9$, $P=0.08$) and S ($t(15)=1.8$, $P=0.09$; all other P s > 0.19). There was no evidence of coding of location in ERC (all F s < 1 , NS). To test whether heading representations significantly differed between RSC and ERC, we submitted the correlations to an ROI x Direction x Similarity repeated measures ANOVA. The overall strength of heading coding did not differ between regions (ROI x Similarity: $F(1,15)=1.4$, $P=0.25$), but the magnitude of coding across directions did differ between RSC and ERC (ROI x Direction x Similarity: $F(3,45)=5.3$, $P=0.003$).

In sum, these analyses characterized the heading and location codes observed in the original ANOVAs. RSC and ERC both coded for heading, but in different ways: RSC only represented N whereas heading coding in ERC did not significantly differ across directions. PPA represented all locations equally well. These analyses also revealed weak spatial effects not identified by the ANOVAs. Specifically, there was marginal evidence for coding of Location 1 in RSC and for coding of N in PPA.

3.5 Discussion

We measured multivoxel activity patterns while subjects performed two versions of a judgment of relative direction (JRD) task that required them to access long-term spatial knowledge of a familiar college campus. By interrogating the cross-task pattern similarities, we have shown that specific regions of temporal and parietal cortex support representations of real-world spatial information that abstract across the nature of the stimulus (picture or text) and the cognitive processes required to access the representation (analysis of the visual scene or mental imagery). Specifically, we found that RSC and ERC represented the initial allocentric heading of the JRD in a way that abstracted across different locations whereas PPA represented the initial location of the JRD in a way that abstracted across different headings.

Our first main result is that both RSC and ERC represented the allocentric heading that was instantiated during the JRDs, a representation that was consistent whether subjects oriented based on a photograph of a real-world scene or imagined

the directional relationship between two landmarks. Previous neuroimaging, neuropsychology, and neurophysiology studies have all provided converging evidence for the importance of RSC in representing allocentric heading. In healthy human subjects, RSC is strongly activated when subjects are asked to retrieve heading information (Epstein et al., 2007), and recent studies have shown that RSC can discriminate allocentric headings within real (Vass and Epstein, 2013) and virtual (Baumann and Mattingley, 2010) environments. When this area of the brain is damaged in humans, it results in “heading disorientation,” an inability to retrieve directional information from environmental stimuli despite preserved knowledge of landmark identities (Aguirre and D'Esposito, 1999). Furthermore, neurophysiological recordings from putative homologous regions in rodents have identified head direction (HD) cells, which fire when the animal's head is oriented in a particular allocentric direction, regardless of the animal's position in the environment (Chen et al., 1994; Cho and Sharp, 2001).

Human and animal studies have also implicated ERC in coding of directional information. Direct recordings of ERC cells from human neurosurgical patients navigating a virtual environment have revealed putative HD cells (Jacobs et al., 2010), and virtual navigation in the fMRI scanner elicits ERC activation characterized by a directional signal with six-fold symmetry, consistent with a population of grid cells (Doeller et al., 2010). Recordings of ERC neurons in rodents have identified grid cells, HD cells, and grid x HD cells in this region (Hafting et al., 2005; Sargolini et al., 2006), all of which modulate their activity on the basis of allocentric heading. Why then were we able to detect ERC heading codes in the current experiment when we were unable to detect them in our previous experiment (Vass and Epstein, 2013)? One possibility is that the heading codes we observe here are supported not by HD cells, but by “path cells,” a recently identified cell type in human ERC which codes for the direction of travel along a path (Jacobs et al., 2010). In the current experiment, three out of four locations are on the same path, Locust Walk, which runs East-West along the length of

campus. Thus, a path cell representation would be characterized by across-location coding of East and West, which is consistent with our observations in ERC. In contrast, the comparisons in Vass & Epstein were more often across-path, which would have diminished our ability to detect such a representation. In any case, when we re-examined the data from our previous experiment, we found that there was marginal evidence of heading codes in Left ERC ($P=0.056$), which is consistent with the current results.

Thus, the current finding of heading codes in RSC and ERC fits well with the existing human and animal spatial cognition literatures. The current results extend these previous findings in two ways. First, they demonstrate that representations of allocentric heading can be detected during mental imagery processes and are not dependent solely on visual input. Second, the representations induced by these top-down processes are consistent with those elicited by bottom-up input (i.e., visual scenes), suggesting that overlapping populations of neurons subserve both functions, a hypothesis put forth by a prominent model of spatial memory (Byrne et al., 2007).

We also found that the nature of the heading representation differed between RSC and ERC: whereas RSC preferentially coded North, the strength of coding in ERC did not significantly differ between directions. Animal studies of navigation have provided a wealth of evidence for ERC-like representations of heading, in which a population of HD cells represents all headings equally (Taube et al., 1990b; Giocomo et al., 2014). To our knowledge, direct recordings from RSC have not indicated an unequal distribution of represented headings among the HD cell population. However, human studies of spatial memory have consistently observed evidence of “orientation-dependent” representations when subjects perform JRD tasks (for a review, see McNamara, 2003). In these studies, subjects are faster and/or more accurate when making judgments aligned to a particular heading (or headings), a result that has been interpreted as indicating that these headings are more accessible in memory (Shelton and McNamara, 1997; Montello et al., 2004). Orientation-dependent spatial memory

has been most commonly studied using small tabletop displays of objects, but has also been reported for large environmental spaces with which subjects have long term navigational experience (Marchette et al., 2011; Frankenstein et al., 2012). We observed behavioral evidence of orientation-dependent representations in both versions of our JRD task, with subjects responding significantly faster for JRDs headed North. RSC showed the same pattern of orientation-dependence and represented North significantly more strongly than the other headings. The preference for North may reflect subjects' previous experience with maps of campus, which are drawn in a North-up fashion, or a more general cultural preference for orienting according to North. Our results are consistent with the hypothesis that even highly familiar environments are represented in memory according to a particular reference direction, and suggest that RSC may be the neural locus of this reference direction representation. In sum, the results from RSC and ERC indicate that there are at least two forms of allocentric heading representation that are activated during the JRD task, one in ERC which represents all headings within the environment (or at least the cardinal directions tested here), and one in RSC which represents the reference direction of the environment. This novel finding is the first evidence in humans of distinct kinds of heading representations during the same task.

We failed to replicate our previous observation of heading codes in presubiculum (Vass and Epstein, 2013). One possible explanation is that the activity patterns associated with a particular heading remapped across tasks or across scan runs. Indeed, neurophysiology studies have shown that remapping is a common occurrence in the spatial memory system after a change in the environment or in the internal state of the navigator. Place cells may exhibit either rate remapping, in which the same ensemble of neurons activate for a particular location but with different firing rates, or global remapping, in which each cell of the ensemble independently changes its place field, or ceases to fire at all (Muller and Kubie, 1987; Leutgeb et al., 2005; Leutgeb and Leutgeb, 2014). Grid cells also exhibit remapping, expressed as a

translation and/or rotation of their firing fields (Fyhn et al., 2007). Moreover, a recent fMRI study measuring the macroscopic grid signal observed a remapping of the orientation of the grids across scan runs, even when the environmental context was held constant (Pape et al., 2011). As such, it would not be surprising if the change in stimuli/task in our experiment induced a remapping of activity patterns. More work will be needed to delineate the conditions under which heading representations remap or remain consistent and how this remapping differs across brain regions.

Our second main result is that PPA coded for the starting location of the JRD problem in a way that was consistent across headings and tasks. In a previous experiment, we observed abstract coding of location in PPA, although this effect was no longer significant when we performed an analysis to remove the variance associated with low-level visual similarity (Vass and Epstein, 2013). Here, where we have explicitly controlled for visual similarity by using two different kinds of stimuli, we find that activity patterns in PPA are consistent across views from the same location, whether those views are perceived or imagined. Furthermore, PPA exhibited equivalent location coding whether locations were defined by punctate objects (N=3; N.B. the landmark object did not appear in the photographs) or a courtyard enclosed by buildings (N=1). What is the mechanism underlying these location codes? Both place cells and spatial view cells have been identified in the parahippocampal gyrus in recordings from human neurosurgical patients (Ekstrom et al., 2003; Jacobs et al., 2010; Miller et al., 2013). Thus, one possibility is that the location coding observed in the current experiment reflects the differential patterns elicited by ensembles of place cells. However, a more likely mechanism is that perceiving or imagining a particular scene activated the representations of nearby views. Indeed, in a recent experiment from our lab, we observed that PPA could cross-decode the interiors of familiar landmarks from their exteriors, even though they have little visual similarity, and that this decoding was abolished when subjects were asked to imagine an object or a face with each scene, thereby disrupting the mental imagery of other scenes (Marchette et al., 2014). In the

current experiment, these nearby views may have been activated by boundary extension, a phenomenon in which scenes are remembered as containing a more expansive view than that which was actually perceived and which PPA has previously been shown to be sensitive to (Park et al., 2007; Chadwick et al., 2013). This interpretation would also be consistent with post-scan debriefings from multiple subjects who reported that their imagined view during the verbal runs was more wide angle than the photographs in the picture runs. An important avenue for future research will be adjudicating between these and other possible mechanisms of location representations in PPA.

In contrast to our previous experiment (Vass and Epstein, 2013), RSC did not exhibit location codes that abstracted across tasks, with the possible exception of weak coding of Location 1. Although there are many experimental design differences between the two experiments, any of which could have contributed to the inability to detect location coding, one particularly important difference may be the tasks that were used. In our previous experiment, subjects were asked to consider one location per trial, but here, subjects were asked to consider either two or three locations per trial in the picture and verbal runs respectively. If RSC represents locations, evoking the representations of multiple locations per trial would have made it more difficult to identify the pattern variance attributable to the starting location. A second possibility is that the nature of the location representation differs across tasks. A previous fMRI study using verbal and picture JRD tasks found that medial parietal and retrosplenial regions were more active during the verbal JRD than the picture JRD (Zhang et al., 2012). Moreover, a previous behavioral experiment observed different task-dependent representations when subjects made two different judgments about the same environment (Valiquette and McNamara, 2007). When subjects performed JRDs, they accessed a single orientation-dependent representation whereas when subjects performed a scene recognition task on the same environment, they accessed two orientation-dependent representations corresponding to the views they had seen

during training. A comparison of the behavioral priming results across our two experiments also supports this interpretation. When subjects made heading judgments in the previous experiment, they exhibited priming if the previous trial was a view of the same location, but not if it was a view facing the same direction, consistent with the hypothesis that they activated a more orientation-independent representation. Here, we observe faster reaction times overall for North-facing trials, and an additional benefit for repetition of North trials in the verbal runs, consistent with the hypothesis that subjects activated a North-oriented representation while performing JRDs. Finally, a third possibility is that the representation in RSC encodes both location and heading information in the two tasks (cf. DiCarlo and Cox, 2007), but the temporal dynamics lead to a location-dominated representation during the heading task, as heading no longer needs to be processed once the solution has been reached, and a heading-dominated representation during the JRD task, as heading needs to be maintained in order to report the correct egocentric target direction. In order to adjudicate between these and other potential explanations, it will be important for future studies to examine the effect of task on the spatial representations elicited in RSC.

In summary, we have demonstrated distributed coding of allocentric spatial information within medial temporal and medial parietal regions when subjects performed complex, ecologically relevant navigation tasks that required them to access the directional relationships between familiar landmarks. Consistent with a prominent model of spatial memory (Byrne et al., 2007), these representations were consistent even across drastic differences in the stimulus that cued spatial memory. We also report the novel observation that RSC and ERC maintained distinct heading codes during the same task, a finding not predicted by previous neurophysiological recordings in animals. Thus, an important future direction will be to further characterize the heading codes in these regions. Does RSC only encode the principal reference direction or is this specific to the JRD task? Does ERC encode all headings, including those not aligned to the geometry of the environment? How do RSC and ERC

representations interact during perception and during spatial memory retrieval?

Addressing these questions will provide a more complete model of human spatial memory and may generate novel predictions that can be interrogated using animal models.

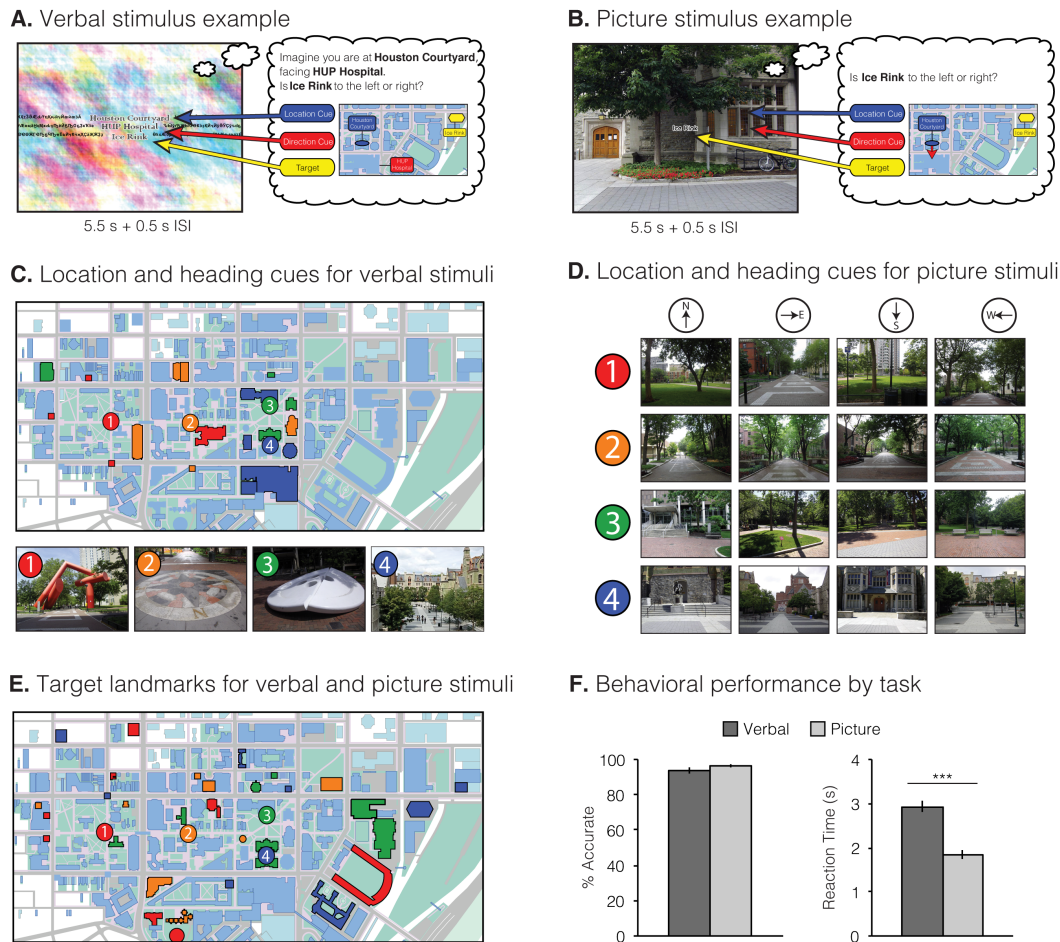



Figure 3.1 Experimental design and task. A. Example stimulus from a verbal scan run. On each trial, subjects indicated whether a target would be on their left or right given an imagined starting location (“Location Cue”) and heading (“Direction Cue”). Text was displayed in 1 of 17 fonts and superimposed on 1 of 17 different colorful backgrounds. B. Example stimulus from a picture scan run. On each trial, subjects indicated whether a target landmark would be on their left or right given the view shown in the photograph, which conveyed both location and heading information. Note that both A & B depict the same JRD, but with different visual cues. C. The same 4 starting locations were used for verbal and picture stimuli (numbered circles; N.B. the landmark object never appeared in the photographs, shown in D). Each of the 4 locations was assigned a unique set of landmarks to serve as North, East, South, and West heading cues. Heading cues are colored according to their respective starting location. D. The picture stimuli consisted of views facing North, East, South, and West at each of the 4 locations. We collected 17 photographs of each view (1 photograph of each view shown). E) Each location was assigned a unique set of 8 target landmarks, which were used in both versions of the task. Each view (N=16) was assigned to one target to the left and one target to the right (32 targets total). F) Subjects were equally accurate on both versions of the task, but were significantly faster for JRDs cued with picture stimuli. *** $P < 0.001$.

A. Priming for location



B. Priming for heading

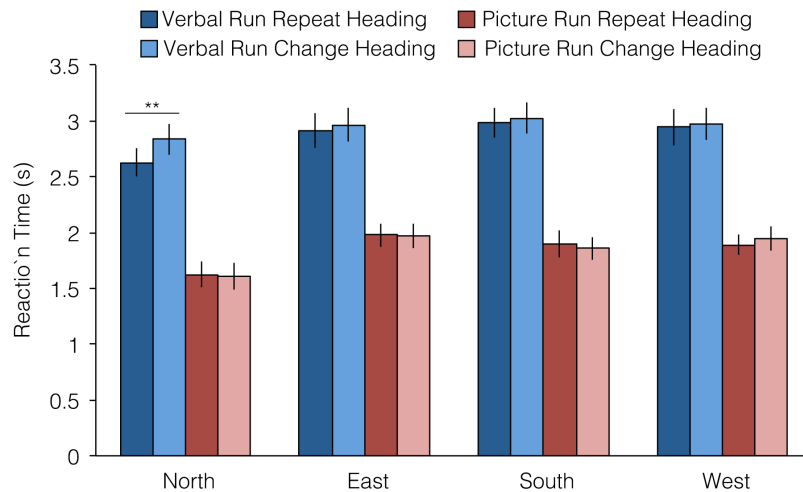
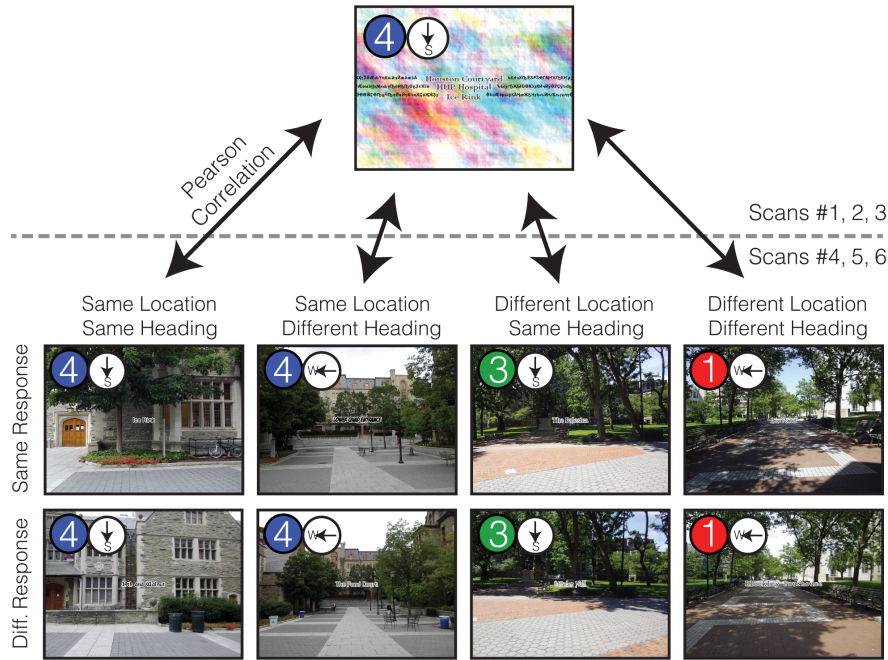


Figure 3.2 Behavioral priming. A. Priming for location. Each trial was sorted based on its starting location and whether the previous trial started at the same location ("Repeat") or a different location ("Change"). In the verbal runs, subjects showed significant priming for all 4 locations, but this condition involves repetition of the words on the screen across trials. There was no location priming in the picture runs. B. Priming for heading. Each trial was sorted based on its heading and whether the previous trial assumed the same heading ("Repeat") or a different heading ("Change"). There was significant heading priming for North, but only in the verbal runs. * $P < 0.05$; ** $P < 0.01$; *** $P < 0.001$

A. Multivoxel pattern correlations



B. Location x Heading x Response ANOVAs




Figure 3.3 Multivoxel pattern correlations in PPA and RSC. A. Coding of abstract location and heading was assessed by measuring the pattern similarities of the JRD trials across verbal and picture runs. We calculated the mean pattern of voxelwise activity across all presentations of each JRD in each version of the task. We then obtained pairwise similarity of all JRDs by calculating the Pearson correlation of the voxelwise activity patterns across verbal and picture runs. JRD pairs were grouped based on whether they shared the same location, heading, and response, creating 8 groups (1 example pair of each shown). B. If a region codes for spatial information, two JRDs that share the same spatial information (e.g., “Same Location”) should be more similar than two JRDs which do not share that spatial information (e.g., “Different Location”). To test for this, we submitted the correlation values for RSC and PPA to a Location x Heading x Response ANOVA. RSC exhibited a main effect of Heading and a Heading x Response interaction. PPA exhibited main effects of Location and Heading and interactions effects of Location x Response and Location x Heading x Response.




Figure 3.4 Multivoxel pattern correlations in MTL ROIs. Location x Heading x Response ANOVAs for the seven MTL ROIs. Inset, bottom left, shows example ROIs from one subject. ERC exhibited a significant main effect of heading. L Presub. exhibited a significant Location x Response interaction. Ant. Hipp., anterior hippocampus; Post. Hipp., posterior hippocampus; L, left; R, right; Presub., presubiculum; ERC, entorhinal cortex; PRC, perirhinal cortex; PHC, parahippocampal cortex.




Figure 3.5 Searchlight analysis for heading. Regions in medial parietal cortex coded for headings in a way that abstracted across locations and tasks. Results are plotted on the inflated surface of one subject's brain, where dark grey represents sulci and light grey represents gyri. Yellow voxels are significant at $P < 0.05$ after correction for multiple comparisons across the entire brain. Outlines display the boundaries of the display subject's PPA and RSC.

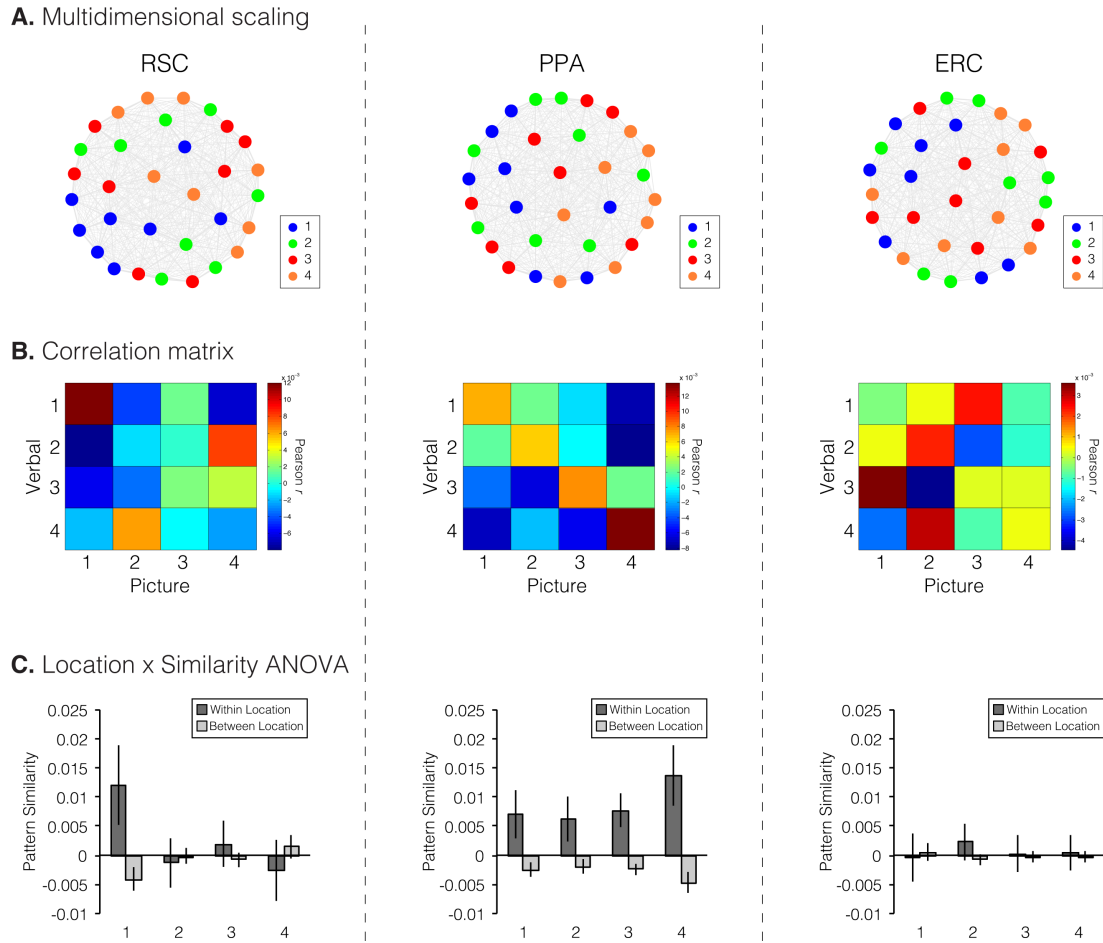




Figure 3.6 Location representation in RSC, PPA, and ERC. A. Multidimensional scaling of JRD pattern similarities for RSC (left), PPA (center) and ERC (right). Each point represents a JRD, and distance between points correspond to correlation distances. Points have been color coded based on the starting location of the JRD. B. Mean cross-task Pearson correlation for all pairs of locations, excluding same-JRD correlations. Rows correspond to locations in the verbal runs and columns correspond to locations in the picture runs. C. Location x Similarity ANOVAs showing mean cross-task within- and between-location correlations for each location. Within-location correlations correspond to elements on the diagonal of the correlation matrix in B and between-location correlations correspond to the average of rows and columns of off-diagonal elements in B. PPA exhibited a significant main effect of location.

A. Multidimensional scaling



B. Correlation matrix



C. Heading x Similarity ANOVA




Figure 3.7 Heading representation in RSC, PPA, and ERC. A. Multidimensional scaling of JRD pattern similarities for RSC (left), PPA (center) and ERC (right). Each point represents a JRD, and distance between points correspond to correlation distances. Points have been color coded based on the heading of the JRD. B. Mean cross-task Pearson correlation for all pairs of headings, excluding same-JRD correlations. Rows correspond to headings in the verbal runs and columns correspond to headings in the picture runs. C. Heading x Similarity ANOVAs showing mean cross-task within- and between-heading correlations for each heading. Within-heading correlations correspond to elements on the diagonal of the correlation matrix in B and between-heading correlations correspond to the average of rows and columns of off-diagonal elements in B. RSC and ERC both exhibited significant main effects of heading, but the heading x similarity interaction was only significant for RSC. N, north; E, east; S, south; W, west.

CHAPTER 4: DISTANCES BETWEEN REAL-WORLD LOCATIONS ARE REPRESENTED IN THE HUMAN HIPPOCAMPUS

Morgan, L.K., MacEvoy, S.P., Aguirre, G.K., and Epstein, R.A. (2011) Distances between real-world locations are represented in the human hippocampus. *Journal of Neuroscience*, 31(4): 1238-45.

4.1 Abstract

Spatial navigation is believed to be guided in part by reference to an internal map of the environment. We used functional magnetic resonance imaging (fMRI) to test for a key aspect of a cognitive map: preservation of real-world distance relationships. University students were scanned while viewing photographs of familiar campus landmarks. fMRI response levels in the left hippocampus corresponded to real-world distances between landmarks shown on successive trials, indicating that this region considered closer landmarks to be more representationally similar and more distant landmarks to be more representationally distinct. In contrast, posterior visually responsive regions such as retrosplenial complex and the parahippocampal place area were sensitive to landmark repetition and encoded landmark identity in their multivoxel activity patterns but did not show a distance-related response. These data suggest the existence of a map-like representation in the human medial temporal lobe that encodes the coordinates of familiar locations in large-scale, real-world environments.

4.2 Introduction

A cognitive map is a representational structure that encodes spatial locations within large-scale, navigable environments. O'Keefe and Nadel (1978) proposed that the hippocampus is the brain structure that supports the cognitive map in mammals. Supporting this hypothesis are data from neurophysiological studies indicating that hippocampal neurons exhibit increased firing for particular spatial locations (O'Keefe and Dostrovsky, 1971; Matsumura et al., 1999) and lesion data indicating that damage to the hippocampus impairs navigation using map-based but not route-based

strategies (Morris et al., 1982). The theory has been further enhanced by the recent discovery of a grid-like spatial representation in entorhinal cortex, the primary source of hippocampal input (Hafting et al., 2005). The spatial regularity of the entorhinal grid suggests that it may facilitate precise coding of location within the environment and a metric for calculating distances between locations (Jeffery and Burgess, 2006).

In humans, the evidence for hippocampal involvement in cognitive map coding is less clear. Although place cells have been discovered in the human hippocampus (Ekstrom et al., 2003), damage to this structure does not lead to a purely spatial impairment. Rather, these amnesic patients suffer from a more general declarative memory problem (Squire, 1992), which can leave the ability to navigate through familiar environments essentially intact (Teng and Squire, 1999). Furthermore, neuroimaging studies of spatial navigation obtain hippocampal activation in some cases (Ghaem et al., 1997; Maguire et al., 1998) but not others (Aguirre et al., 1996; Aguirre and D'Esposito, 1997; Rosenbaum et al., 2004). In summary, the claim that human medial temporal lobe structures such as hippocampus encode spatial information per se, as opposed to other kinds of navigationally relevant information, remains controversial (Shrager et al., 2008).

Here we present evidence for a signal in the human hippocampus that exhibits a key feature of a cognitive map: preservation of real-world distance relationships. That is, the hippocampus considers locations that are physically closer in space to be more representationally similar and locations that are further apart in space to be more representationally distinct. Such a distance-related response has not been identified previously in the hippocampus: the existence of place cells indicates that different locations are distinguished but does not necessarily imply that these locations are organized according to a map-like code. To test for such a code, we scanned university students with functional magnetic resonance imaging (fMRI) while they viewed images of landmarks from a familiar college campus. We examined multivoxel activity patterns evoked by landmarks as well as adaptation effects related to the

distance between landmarks. We reasoned that a brain region involved in encoding locations within an allocentric map should demonstrate adaptation effects that are proportional to the real-world distance between successively viewed landmarks. In contrast, regions representing visual or semantic information about landmarks should exhibit adaptation during landmark repetition and multivoxel patterns that distinguish between landmarks but should not exhibit distance- related adaptation.

4.3 Materials and Methods

SUBJECTS. Fifteen right-handed volunteers (10 female; mean age, 22.6 ± 0.3 years) with normal or corrected-to-normal vision were recruited from the University of Pennsylvania. All subjects had at least 1 year of experience with the campus (average length of experience, 3.7 ± 0.2 years) and gave written informed consent according to procedures approved by the University of Pennsylvania institutional review board.

MRI ACQUISITION. Scans were performed at the Hospital of the University of Pennsylvania on a 3 T Siemens Trio scanner equipped with a Siemens body coil and an eight-channel head coil. High-resolution T1- weighted anatomical images were acquired using a three-dimensional magnetization-prepared rapid-acquisition gradient echo pulse sequence [repetition time (TR), 1620 ms; echo time (TE), 3 ms; inversion time (TI), 950 ms; voxel size, $0.9766 \times 0.9766 \times 1$ mm; matrix size, $192 \times 256 \times 160$]. T2*-weighted images sensitive to blood oxygenation level- dependent contrasts were acquired using a gradient-echo echo-planar pulse sequence (TR, 3000 ms; TE, 30 ms; voxel size, $3 \times 3 \times 3$ mm; matrix size, $64 \times 64 \times 45$). Images were rear-projected onto a Mylar screen at 1024×768 pixel resolution with an Epson 8100 3-LCD projector equipped with a Buhl long-throw lens. Subjects viewed the images through a mirror attached to the head coil. Images subtended a visual angle of $22.9^\circ \times 17.4^\circ$.

STIMULI AND PROCEDURE. Visual stimuli were color photographs of 10 prominent landmarks (i.e., buildings and statues) from the University of Pennsylvania campus. Twenty-two distinct photographs were taken of each landmark for a total of 220 images. To ensure that all subjects were familiar with the landmarks, they underwent behavioral testing 1 d before scanning in which they were asked to indicate (yes/no) whether they were familiar with each landmark. In the same session, “subjective” distances between landmarks were determined by asking subjects to estimate the number of minutes required to walk between each pair of locations.

The main experiment consisted of two fMRI scan runs that lasted 6 m 51 s each, during which subjects viewed all 220 images without repetition. Images were presented every 3 s in a continuous-carryover sequence that included 6 s null trials interspersed with the stimulus trials (Aguirre, 2007). This stimulus sequence counterbalances main effects and first- order carryover effects, thus allowing us to use the same fMRI dataset to examine both the multivoxel response pattern for each landmark and adaptation between landmarks presented on successive trials. A unique continuous-carryover sequence was defined for each subject. On each stimulus trial, an image of a landmark was presented for 1 s, followed by 2 s of a gray screen with a black fixation cross. Subjects were asked to covertly identify each campus landmark and make a button press once they had done so. During null trials, a gray screen with black fixation cross was presented for 6 s during which subjects made no response. Each run included a 15 s fixation period at the beginning of the scan to allow tissue to reach steady-state magnetization and ended with an additional 15 s fixation period.

After the experimental runs, subjects were scanned twice more for the functional localizer. Each functional localizer scan lasted 7 m 48 s and consisted of 18 s blocks of images of places (e.g., cityscapes, landscapes), single objects without backgrounds, scrambled objects, and other stimuli, presented for 490 ms with a 490 ms interstimulus interval.

DATA PREPROCESSING. Functional images were corrected for differences in slice timing by resampling slices in time to match the first slice of each volume, realigned to the first image of the scan, and spatially normalized to the Montreal Neurological Institute (MNI) template. Data for all univariate analyses, including the functional localizer scans, were spatially smoothed with a 6 mm full-width half-maximum (FWHM) Gaussian filter; data for multivoxel pattern analyses (MVPAs) were left unsmoothed.

REGIONS OF INTEREST. Data from the functional localizer scans were used to define functional regions of interest (ROIs) for scene-responsive cortex in parahippocampal place area (PPA) and retrosplenial complex (RSC) (places > objects), object-responsive cortex in the lateral occipital complex (LOC) (objects > scrambled objects), and early visual areas (scrambled objects > objects). Thresholds were determined on a subject-by-subject basis to be consistent with those identified in previous studies and ranged from $T > 2.0$ to $T > 3.5$ (mean $T = 2.7 \pm 0.1$). Bilateral PPA and LOC were located in all 15 subjects. Right RSC was identified in all subjects and left RSC in 13 of 15 subjects. We also defined anatomical ROIs for the hippocampus using sagittal T1-weighted images. The hippocampal ROI included all CA fields and the subiculum but did not include entorhinal cortex. The hippocampus was separately defined for the left and right hemispheres and further subdivided into its anterior/ inferior and posterior/superior subregions by an axial division at $z = -9$.

UNIVARIATE ANALYSES. Data were analyzed using the general linear model as implemented in VoxBo (www.voxbo.org), including an empirically derived 1/f noise model, filters that removed high and low temporal frequencies, and nuisance regressors to account for global signal variations and between-scan differences. Between-landmark adaptation effects were modeled with a regressor corresponding to the distance between each landmark and the immediately preceding landmark, calculated in one of two ways: the Euclidean distance in meters between landmarks

(i.e., objective distance “as the crow flies”) or an individual subject’s perceived distance in minutes of travel time between landmarks (subjective distance), each mean centered. Also included in the model was a regressor modeling the response to any landmark versus baseline and two regressors to account for situations in which the distance to the previous landmark was undefined: (1) when a landmark stimulus followed a null trial and (2) when a stimulus consisted of a landmark after another view of the same landmark (i.e., repeated-landmark trials). A separate, supplementary analysis examined distance effects in a less constrained manner by assigning each trial to one of four bins based on the distance from the currently presented to the previously presented landmark, plus a fifth regressor for repeated-landmark trials. Finally, a modified version of the first model was run, in which distance was only defined for non-covisible landmarks, and the covisible versus non-covisible distinction was modeled with an additional regressor.

For all models, β values were calculated for each ROI, which were then compared with zero using a one-tailed t test. In addition, whole-brain analyses were performed by calculating subject-specific t maps for contrasts of interest, which were then entered into a second-level random-effects analysis. Monte Carlo simulations involving sign permutations of the whole-brain data from individual subjects (1000 relabelings, 12 mm FWHM pseudo- t smoothing) were performed to find the true type I error rate for each contrast (Nichols and Holmes, 2002). All reported voxels are significant at $p < 0.05$, corrected for multiple comparisons across the entire brain. To ensure accurate localization of distance-related adaptation effects to the hippocampus in the whole-brain analyses, we performed an additional step to anatomically coregister the structures of the medial and lateral temporal lobes for this contrast. The hippocampus, entorhinal cortex, perirhinal cortex, parahippocampal cortex, insula, superior temporal gyrus, and middle temporal gyrus were anatomically defined according to parcellation protocols (Kim et al., 2000; Matsumoto et al., 2001; Pruessner et al., 2002; Kasai et al., 2003). These structures were then coregistered

across subjects using the ROI alignment method and the same transformations applied to the functional data before random-effects analysis (Yassa and Stark, 2009). The results were similar when this additional coregistration step was not performed.

MULTIVARIATE ANALYSES. Twenty regressors were created to model each of the 10 landmarks separately within the two experimental runs. These regressors were then used to extract β values for each condition at each voxel. Multivoxel pattern classification was performed on these values using custom MATLAB code based on the method described by Haxby et al. (2001). In short, a cocktail mean pattern was calculated for each of the two runs and subtracted from each of the individual patterns before classification. Pattern classification was performed by pairwise comparisons across all 10 landmarks. Patterns were considered correctly classified if the average pattern correlation between landmark A in opposite halves of the data was higher than between landmark A and landmark B in opposite halves of the data. Classification accuracy was then averaged across all possible pairwise comparisons for a given ROI and tested against random chance (i.e., 0.5) using a one-tailed t test. We also examined classification using a one-versus-all procedure in which landmark A was only considered correctly classified if the same-landmark correlation between opposite halves of the data (i.e., landmark A–landmark A) was higher than all nine cross-landmark correlations (i.e., landmark A–landmark B, landmark A–landmark C, etc.). Chance in this analysis is 10%.

A searchlight analysis based on Kriegeskorte et al. (2006) was implemented using custom MATLAB code to look for areas of high classification accuracy outside of the predefined ROIs. A small spherical ROI (radius, 5 mm) was created and centered on each voxel of the brain in turn. Overall classification accuracy was calculated for this region using the pairwise comparison procedure, and the value was assigned to the center voxel of the cluster. These values were then used to create subject-specific accuracy maps, which were smoothed with a 9 mm FWHM Gaussian kernel before

entry into a random-effects analysis. As before, a Monte Carlo sign permutation test was performed to calculate the true false-positive rate for classification accuracy against chance (50%). All reported voxels are significant at $p < 0.05$, corrected for multiple comparisons across the entire brain. To test whether landmarks that are nearer in space have more similar multivoxel patterns, we computed correlations between neural distance and physical distance. Neural distance between two landmarks A and B was quantified as $1 - r_{AB}$, where r_{AB} is the correlation between the pattern elicited by landmark A and the pattern elicited by landmark B after subtraction of the cocktail mean from both. Because this analysis does not require reserving part of the fMRI data as a separate test set, the fMRI response patterns used in this calculation included data from both scan runs. Neural distances were obtained for all pairs of landmarks and were correlated with the actual physical distances between those pairs. Pearson's R values were then converted to Fisher's Z values, averaged across subjects, and compared against zero using a one-tailed t test. This analysis was performed both within predefined ROIs and also within a set of 5 mm searchlights whose center positions covered the entire brain.

4.4 Results

BEHAVIORAL RESPONSES. During the main experiment, University of Pennsylvania students viewed photographs of prominent landmarks (buildings and statues) from the Penn campus (Figure 4.1), which were presented one at a time without any image repetitions. Subjects made a button press once they identified the landmark shown on each trial. Note that this task did not explicitly require subjects to retrieve information about the location of the landmark or its relationship to other landmarks. Reaction times on this task revealed a behavioral priming effect for landmark identity: subjects responded more quickly on trials in which the landmark was a repeat of the landmark shown on the previous trial than on non-repeat trials (repeat, 522 ± 29 ms vs. nonrepeat, 547 ± 30 ms; $t(14) = -2.0$, $p = 0.03$). We also measured reaction time as a

function of the real-world distance between the currently viewed landmark and the landmark shown on the previous trial; however, here we observed no significant effect ($r = 0.002$, $p = 0.48$).

FMRI ADAPTATION ANALYSES. fMRI adaptation is a reduction in response observed when an item is repeated, or when elements of an item are repeated (Grill-Spector et al., 2006). This reduction is interpreted as indicating representational overlap between the first and second item, with the amount of adaptation proportional to the degree of overlap (Kourtzi and Kanwisher, 2001). We examined two forms of fMRI adaptation effects within our functionally and anatomically defined ROIs. First, we looked for adaptation effects caused by presentation of the same landmark on successive trials. When the landmark on the current trial was identical to the landmark shown on the preceding trial, fMRI responses in PPA and RSC were significantly attenuated, as indicated by a significant negative loading on a regressor modeling response differences between repeat and nonrepeat trials (PPA, $t(14) = -3.25$, $p = 0.003$; RSC, $t(14) = -3.47$, $p = 0.002$). Whole-brain random-effects analysis revealed additional landmark-related adaptation in the left superior lingual gyrus abutting the anterior calcarine sulcus (-18, -53, 1) and the left medial retrosplenial region (-6, -47, 15) medial to the functionally defined RSC (Figure 4.2). At lower thresholds, these activations extended into the functionally defined RSC and the PPA/fusiform region.

Next, we looked for adaptation between pairs of landmarks as a function of the real-world distance (i.e., objective distance) between them. We predicted that regions supporting a map-like representation would exhibit greater adaptation (i.e., less fMRI response) when proximal landmarks were shown on successive trials and less adaptation (i.e., greater fMRI response) when distal landmarks were shown on successive trials. We tested for a linear relationship between neural response and the distance between the currently viewed landmark and the landmark shown on the immediately preceding trial by measuring the loading on a continuous covariate

modeling real-world distances between successive trials. This effect was positive and significant in the left anterior hippocampus ($t(14) = 4.35$, $p = 0.0003$), indicating that activity in this region correlated with real-world distances between sequentially presented landmarks. This effect was confined to the left anterior hippocampus: no similar relationship was observed in the left posterior ($t(14) = 0.20$, $p = 0.42$), right anterior ($t(14) = 0.21$, $p = 0.42$), or right posterior ($t(14) = 0.49$, $p = 0.32$) hippocampal subregions. An analysis of second-order distance (i.e., distance between the current landmark and the landmark occurring two trials back) found no significant effects in any hippocampal subregion (all p values > 0.3).

Because a cognitive map of the environment may not be entirely faithful to the real world, we also assessed the relationship between adaptation effects and subjects' perceived "subjective" distance between landmarks. Subjective distances were estimates of the number of minutes required to walk between each pair of locations, obtained the day before the fMRI scan in a separate testing session. Subjective distance judgments were highly correlated with objective physical distances (mean $r = 0.90$, $p = 1.71 \times 10^{-13}$), as one would expect given the high degree of familiarity with the campus and the grid-like organization of campus paths that facilitate direct or near-direct travel between locations. We found that activation was dependent on subjective distance in the left anterior hippocampus ($t(14) = 3.22$, $p = 0.003$) but no other hippocampal subregions (left posterior, $p = 0.47$; right anterior, $p = 0.47$; right posterior, $p = 0.17$).

Whole-brain analyses revealed significant dependence of activation on objective distance in the left anterior hippocampus ($-29, -9, -18$), consistent with the ROI analyses reported above (Figure 4.3A). Distance-related activation was also observed in the left inferior insula ($-45, -1, -6$ and $-42, -15, -6$), left anterior superior temporal sulcus (aSTS) ($-48, -6, -18$), and right posterior inferior temporal sulcus (pITS) ($46, -62, -2$) near the location usually occupied by middle temporal/medial superior

temporal visual areas (MT/MST) (Kourtzi et al., 2002) (Figure 4.3A). Whole-brain analyses using subjective distances were similar.

To further explore the distance-related adaptation effect in the hippocampus, we performed two additional analyses. First, we passed functional data to a model in which distances between landmarks on successive trials were discretized into four covariates. This allowed us to graphically examine activation as a function of distance without assuming a linear relationship. The results confirm our previous findings (Figure 4.3 B, C) indicating that activity in the left anterior hippocampus scales with distance between campus locations. Second, we performed an analysis in which successively presented landmarks that are covisible (i.e., one landmark can be seen from the other landmark) were modeled separately from landmarks that are not covisible. Distance-related adaptation was then examined for the non-covisible landmarks (because there was little variability in distance for the covisible landmarks). We observed greater activity in the left anterior hippocampus for non-covisible landmarks compared with covisible landmarks ($t(14) = 2.49$, $p = 0.01$), as well as distance-related adaptation among the non-covisible landmarks ($t(14) = 2.97$, $p = 0.005$). This last effect is of particular importance because it indicates that the adaptation effect we have observed cannot be solely attributed to adaptation for landmarks that sometimes occur within the same scene but rather reflects a true distance effect.

Finally, we tested whether distance-related adaptation was found in the regions showing landmark-specific adaptation in the whole-brain analysis and whether landmark-specific adaptation could be found in the regions showing a distance-related effect. We observed a complete dissociation: there was no effect of landmark repetition in the regions showing distance-related adaptation [left anterior hippocampus ($t(14) = -0.20$, $p = 0.42$), left inferior insula ($t(14) = -0.76$, $p = 0.23$), left aSTS ($t(14) = -0.68$, $p = 0.25$), and right pITS ($t(14) = 1.38$, $p = 0.09$)], and there was no effect of distance in the regions sensitive to landmark repetition [superior lingual ($t(14) = -0.86$, $p = 0.20$) and retrosplenial ($t(14) = 0.47$, $p = 0.32$)]. To confirm the apparent dissociation between

brain regions, we performed an analysis (distance, landmark repetition) x ROI ANOVA for three ROI pairings: hippocampus–PPA, hippocampus–lingual gyrus, and hippocampus–retrosplenial cortex. The interaction term was significant for all three pairings [hippocampus–PPA ($F(1,14) = 7.78$, $p = 0.01$), hippocampus–lingual ($F(1,14) = 17.58$, $p = 0.001$), and hippocampus–retrosplenial ($F(1,14) = 13.64$, $p = 0.002$)]. The fact that we did not observe landmark-specific adaptation in the hippocampus although we observed distance-related adaptation may at first seem surprising, but it is in fact similar to findings from other studies indicating that same-identity repetitions engage additional processes not engaged by different-identity repetitions (Sternberg, 1998; Drucker and Aguirre, 2009). Landmark repetition trials were relatively rare in our experiment, and this fact may have led to the engagement of novelty or oddball processing mechanisms on these trials that would have masked or attenuated any adaptation effect (Strange and Dolan, 2001; Summerfield et al., 2008).

MULTIVOXEL PATTERN ANALYSES. A second method for determining the representational distinctions made by a brain region is to examine multivoxel patterns elicited by different stimuli. MVPA can provide information that is complementary to that obtained through adaptation, insofar as MVPA is likely to be more sensitive to information coded on a coarser spatial scale (Drucker and Aguirre, 2009). We performed two such analyses: the first examining the distinguishability of patterns elicited by the 10 campus landmarks, the second examining whether the similarities between these patterns reflected real-world distances.

We first used MVPA to decode the identities of campus landmarks viewed in one scan from patterns evoked during the other scan. This analysis involved comparison of same-landmark and different-landmark patterns across all landmark pairs. Decoding accuracy was significantly above chance in a variety of visually responsive regions (Figure 4.4), including the PPA ($t(14) = 6.12$, $p = 0.00001$), RSC ($t(14) = 4.47$, $p = 0.0003$), object-selective LOC ($t(14) = 7.28$, $p = 0.000002$), and early

visual cortex ($t(14) = 5.18$, $p = 0.00009$). Performance was not significantly different from chance in any of the hippocampal subregions (left anterior, $t(14) = 0.07$, $p = 0.47$; left posterior, $t(14) = 0.77$, $p = 0.23$; right anterior, $t(14) = -0.04$, $p = 0.49$; right posterior, $t(14) = -0.88$, $p = 0.20$). Similar levels of significance were observed when classification performance was scored using a one-versus-all rather than a pairwise comparison procedure. Classification using this method was significantly above chance (10%) in PPA (19.2%, $p = 0.001$), RSC (14.2%, $p = 0.03$), LOC (21.3%, $p = 0.00002$), and early visual cortex (23.6%, $p = 0.0003$) but at chance in the left anterior hippocampus (11.3%, $p = 0.23$). A separate analysis of pairwise decoding performance for individual landmarks indicated that classification performance was approximately equivalent for all landmarks in PPA, RSC, LOC, and early visual cortex and equivalently at chance in the hippocampus (Figure 4.7). This suggests that above-chance classification accuracy is not driven by high performance on only a few landmarks.

A searchlight analysis of pairwise decoding performance across the entire brain revealed areas throughout the occipital and parietal cortices in which landmark identity could be decoded at rates that were significantly above chance (Figure 4.5). Interestingly, these regions were only partially overlapping with regions showing landmark-related adaptation effects in the previous analysis. Similar disjunctions between regions exhibiting adaptation for a stimulus dimension and regions exhibiting multivoxel patterns that distinguish between items along this dimension have been reported previously in the literature (Drucker and Aguirre, 2009).

A second set of analyses tested whether similarities and differences between the multivoxel patterns evoked by the various landmarks related to the real-world distances between the landmarks. To examine this possibility, we calculated a “neural distance” between landmarks for all landmark pairs and then compared this neural distance with the physical distance between landmarks (see Materials and Methods). There was no significant correlation between neural and physical distance in the left anterior hippocampus (mean $r = 0.02$, $p = 0.23$) or in any of the other three

hippocampal subregions (left posterior, mean $r = 0.01$, $p = 0.40$; right anterior, mean $r = -0.02$, $p = 0.28$; right posterior, mean $r = 0.04$, $p = 0.07$). We also examined the correlation between neural and physical distance in the three extrahippocampal regions that exhibited distance-related adaptation. This relationship was not significant in the left aSTS (mean $r = -0.02$, $p = 0.32$), but there was a nonsignificant trend in the right pITS region (mean $r = 0.09$, $p = 0.06$) and a small reversed effect in the left inferior insula (mean $r = -0.06$, $P = 0.02$). A searchlight analysis examining the neural versus physical distance relationship across the entire brain found no significant voxels at either a corrected ($p < 0.05$) or uncorrected ($p < 0.001$) significance level. Levels of performance within the predefined ROIs were not significantly improved by a two-step procedure in which data from one scan run were used for feature selection through a searchlight procedure and testing was performed within the best-performing searchlight on the data from the other scan run (Chadwick et al., 2010).

SUBJECTIVE REPORTS. To gain insight into the cognitive processes that might be driving our observed neural effects, we examined an additional 10 subjects in a purely behavioral version of the experiment, after which they were queried about the thoughts and mental processes they experienced while viewing the campus photographs. This version of the experiment was identical to the fMRI version, except that stimuli were presented on a desktop computer screen within a quiet room. Most subjects (9 of 10) reported they visualized themselves standing at the location the photograph was taken (e.g., “I see Huntsman [Hall] all the time because I’m always in class there, so I was just picturing myself looking at it from this point of view”). Some subjects (6 of 10) noted that the photographs elicited specific memories tied to the viewed locations. For example, one subject reported that a picture taken underneath a campus bridge reminded them of a time when they had walked under it to avoid seeing someone, whereas another subject reported that photographs of the athletic field reminded him of attending a music festival at that location. Only a minority of subjects (3 of 10) reported that they

imagined traveling between the locations. These results suggest that subjects experienced vivid retrieval of the corresponding campus location when viewing the landmark photographs but did not typically have explicit retrieval of the spatial relationships between these landmarks.

4.5 Discussion

DISTANCE-RELATED CODING. Our results demonstrate that fMRI activity in the human hippocampus is modulated by distances between locations in a spatially extended environment. When subjects viewed images of landmarks drawn from a familiar university campus, hippocampal response to each landmark was dependent on the distance between that landmark and the landmark shown on the preceding trial. We observed this distance-related effect although subjects were not given any explicit navigational task but were simply asked to think about the identity of each landmark, suggesting that the mechanism operates essentially automatically. These data are broadly consistent with the idea that the hippocampus either supports a spatial map of the environment or receives direct input from such a map.

These findings advance our understanding of the role of the human medial temporal lobe in spatial navigation. Although previous neuroimaging studies have obtained activation in the hippocampus during virtual navigation and spatial learning (Ghaem et al., 1997; Maguire et al., 1998; Shelton and Gabrieli, 2002; Wolbers and Buchel, 2005; Spiers and Maguire, 2006; Suthana et al., 2009; Brown et al., 2010), this finding is by no means universal (Aguirre et al., 1996; Aguirre and D'Esposito, 1997; Rosenbaum et al., 2004). More importantly, although these studies generally implicated the hippocampus in navigation-related processing, they did not demonstrate hippocampal coding of spatial information per se. A true spatial code does not merely distinguish between different locations (e.g., place A is different from place B) but also encodes the coordinates of those locations such that distance relationships can be

ascertained (e.g., A is closer to B than to C). It is such a distance-preserving code that we demonstrate for the first time here.

Distance-related adaptation effects were also observed in the insula, aSTS, and pITS. Because these effects were unexpected, we interpret them with some caution. Nevertheless, it is intriguing that the pITS region is near the coordinates typically reported for visual areas MT/MST and also exhibited a relationship between interlandmark distance and neural distance for multivoxel patterns. MT/MST has been implicated in the coding of location during virtual navigation tasks such as triangle completion (Wolbers et al., 2007), and neurons with place-selective responses have been observed in this region in monkeys (Froehler and Duffy, 2002). These results suggest that the role of MT/MST in coding location-based information deserves more attention. The insula has also been activated in previous studies of navigation and has been associated with imagined body movements, although its exact role in navigational processing is unknown (Ghaem et al., 1997; Hartley et al., 2003).

In contrast to the adaptation results, similarities between multivoxel patterns in the left anterior hippocampus did not relate to real-world distances between locations. Previous work suggests that multivoxel patterns may be more sensitive to information coded by narrowly tuned neurons clustered by their response properties, whereas adaptation is more sensitive to information coded by broadly tuned neurons with no clustering principle (Drucker and Aguirre, 2009). Thus, finding adaptation effects in the hippocampus but no correlation between distributed patterns and real-world distances suggests a population of neurons with broadly tuned place fields and little spatiotopic organization (Redish et al., 2001). Alternatively, it is possible that the spatial resolution of our study was insufficient for revealing multivoxel patterns in the hippocampus. Using smaller voxels than those used here, a recent study was able to decode the locations of subjects within a virtual-reality room based on hippocampal multivoxel patterns (Hassabis et al., 2009). Although some of the discrepancy between those results and our own may reflect task and analysis differences, it is also possible that location

information would have been evident in the current experiment had the fMRI data been acquired at a finer resolution.

LANDMARK-RELATED CODING. Complementary to the distance-related adaptation effects observed in the hippocampus, landmark-specific adaptation effects were observed in neocortical regions, including the superior lingual gyrus, medial retrosplenial cortex, and (at lower thresholds) RSC and PPA. Our findings are broadly consistent with previous work that indicated these regions code individual scenes and landmarks, but there are two important differences. First, we observed repetition effects in the PPA and RSC, although exact landmark views were never repeated. Thus, the adaptation effect exhibited some degree of viewpoint tolerance. We previously observed cross-viewpoint adaptation in the PPA and RSC when campus scenes were repeated across intervals of several minutes but viewpoint-specific adaptation for shorter repetitions of 100 – 700 ms (Epstein et al., 2008). The present results suggest that intermediate repetition intervals of 2 s elicit viewpoint-tolerant responses more consistent with the longer-interval repetition regimen, a surprising finding that may have important implications for our understanding of the mechanisms that drive fMRI adaptation. Second, previous studies revealed repetition effects primarily in the PPA and RSC, whereas the strongest effects in the current study were found in the medial retrosplenial region abutting, but distinct from, the functionally defined RSC. This region, corresponding to anatomically defined retrosplenial cortex (i.e., Brodmann's areas 29 and 30), has been shown previously to contain spatial and episodic memory-related signals (Rosenbaum et al., 2004; Vann et al., 2009). Thus, the current results emphasize the importance of this region in the retrieval of information about familiar places.

We also examined the multivoxel patterns associated with different campus landmarks. Landmark identity could be decoded in several cortical regions, including some involved in scene perception (PPA, RSC), some involved in object recognition

(LOC), and early visual cortex. These results extend previous findings indicating multivoxel patterns in these regions contain information about scene category (Walther et al., 2009) by showing that they also contain information about specific landmarks. Because all of the stimuli in the current experiment were outdoor images of a college campus, it is unlikely that landmark decoding reflects categorical differences. Rather, these regions may encode visual or geometric properties that are useful for discriminating scenes in terms of general scene categories or as specific scene exemplars. Although these properties may be more holistic in regions such as PPA and RSC, it is likely that simpler visual features such as texture or color may give rise to successful decoding in early visual cortex. In any case, the MVPA and adaptation results converge to implicate neocortical regions such as the PPA and RSC in landmark identification, a role that contrasts with medial temporal lobe involvement in calculating distances between landmarks.

MECHANISMS AND IMPLICATIONS. What are the mechanisms underlying the distance-related signal? The simplest account is that it reflects adaptation among neurons with large and partially overlapping place fields. However, simple adaptation effects in the hippocampus are rarely reported (Brown et al., 1987); thus, we favor an account in which these effects are interpreted in terms of the operation of an active mechanism.

One possibility is that hippocampal activity reflects replay of the route from the immediately preceding landmark to the currently viewed landmark, an operation that would involve more extensive processing for longer routes (Foster and Wilson, 2006). However, we think such an account is unlikely because the subjects did not actually navigate between locations, nor did they report mentally doing so.

Another possibility is that the hippocampal signal reflects the operation of a “mismatch” mechanism that occurs subsequent to an initial pattern completion phase (Gray and McNaughton, 1982; Vinogradova, 2001; Kumaran and Maguire, 2007).

Previous studies have demonstrated that the left hippocampus (but not the right) activates when the expectations of a previously established “context” are violated: for example, when the first few items of a sequence are presented in a familiar order but the last few items are rearranged (Kumaran and Maguire, 2006). In the current experiment, viewing a familiar landmark may have established a “context” on each trial; the hippocampal response on the immediately subsequent trial might then reflect the degree to which the new landmark violated this context. If the activated context on each trial included information about the spatial location of the landmark (in addition, possibly, to nonspatial information not tested here), then the degree of “mismatch” would scale with the distance between landmarks. Alternatively, the degree of context violation might reflect overlap in routes emanating from the two locations, a possibility we cannot exclude given that route overlap is likely to be highly correlated with Euclidean distance on the Penn campus.

Under this account, the hippocampus may work in concert with other brain regions to form a cognitive map. Indeed, based on the rodent data (Hafting et al., 2005) and recent neuroimaging results (Doeller et al., 2010), we suggest that the entorhinal cortex encodes metric information about the spatial relationships between landmarks, whereas the hippocampus calculates the extent to which the current stimulus is consistent or inconsistent with these spatial relationships. This hippocampal–entorhinal representation of the enduring spatial structure of the environment might project to goal representations in the subiculum or other areas, allowing the system to construct routes to different goal locations during navigation (Burgess et al., 2000). Consistent with this hypothesis, Spiers and Maguire (Spiers and Maguire, 2007) observed activity in the subiculum and entorhinal cortex corresponding to distance to a navigational goal; here we show that a different medial temporal lobe region (the anterior hippocampus) encodes distances between landmarks even in the absence of a navigational goal.

The current results may help to illuminate some of the apparent discrepancies between rodent and human data on hippocampal function. Neurophysiological data

(mostly from rodents) indicate that the hippocampus primarily [but not exclusively (Leutgeb et al., 2005; Manns and Eichenbaum, 2009)] encodes spatial information, whereas neuropsychological data (mostly from humans) suggest that hippocampal damage leads primarily to impairments in episodic memory. The idea of context has been used to bridge the gap; indeed, behavioral data indicate that spatial context may play a privileged role in shaping episodic memory (Nadel and Willner, 1980; Hupbach et al., 2008). In the current study, subjects did not physically or mentally navigate between landmarks, but the hippocampal response indicated sensitivity to the spatial relationships between landmarks. We believe that this response may reflect the operation of a spatial context processing mechanism that automatically shapes episodic memory encoding and retrieval.

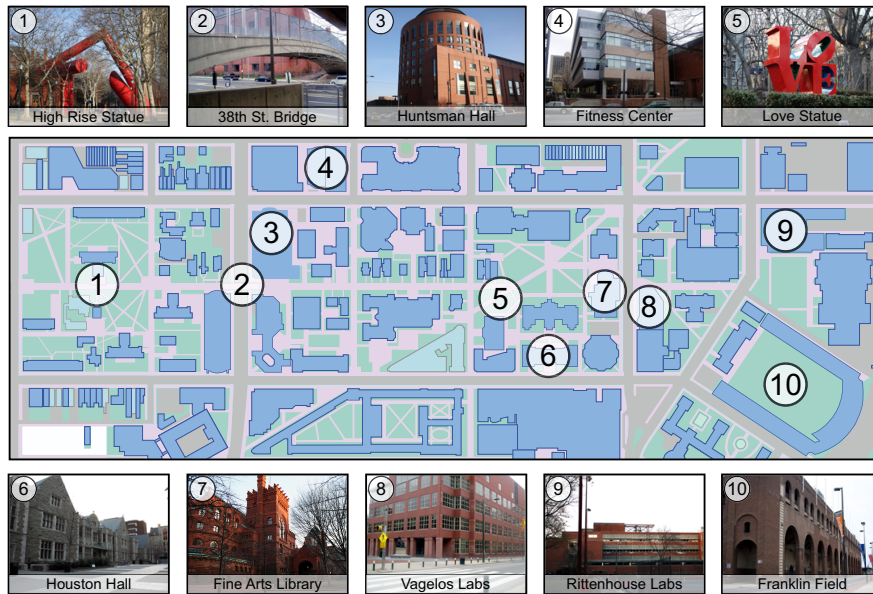


Figure 4.1 Examples of stimuli and map showing the locations of the 10 landmarks on the University of Pennsylvania campus. Twenty-two distinct photographs were taken of each landmark. For more stimulus examples, see Figure 4.6.




Figure 4.2 Whole-brain analysis for landmark adaptation. Voxels showing significant response attenuation when the same landmark was viewed on successive trials are plotted on coronal slices of the MNI template brain. Landmark repetition led to reduced fMRI response in the left superior lingual gyrus (ling) and left medial retrosplenial (retrospl) regions. Landmark-related adaptation was also observed in the PPA and RSC at lower significance thresholds.

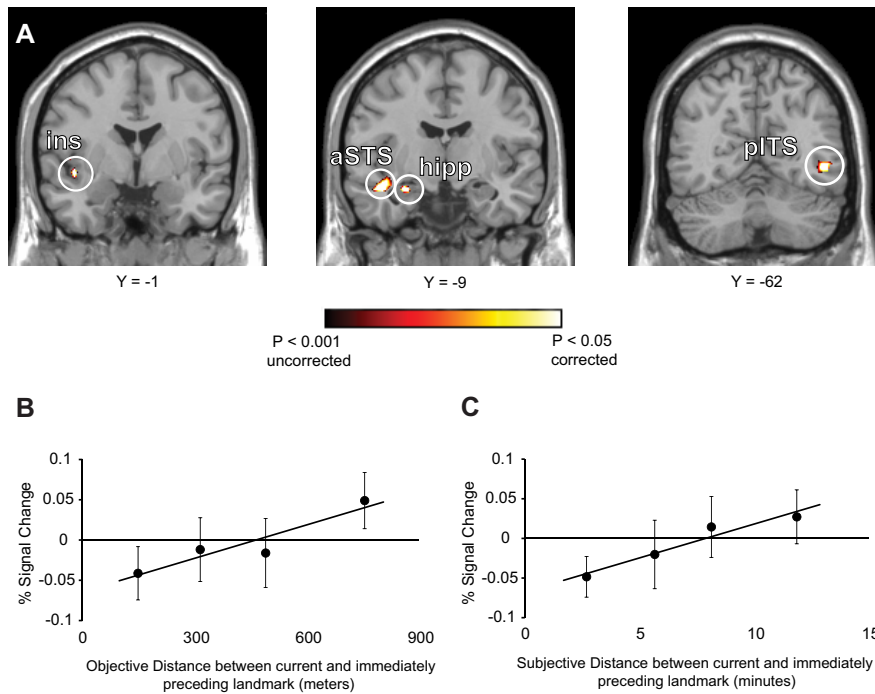


Figure 4.3 Distance-related adaptation in the human brain. A) Colored voxels exhibit fMRI response that scales linearly with real-world distances between landmarks shown on successive trials. Distance-related adaptation was observed in the left inferior insula (ins), left aSTS, left anterior hippocampus (hipp), and right pITS. B) fMRI response (mean \pm SEM percentage signal change) in the anatomically defined left anterior hippocampus plotted as a function of the real-world distance between successively presented landmarks. C) The same plot for subjective distance. fMRI response in the left anterior hippocampus to repeated-landmark (0-distance) trials was 0.016, which was not significantly different from zero ($t = 0.23$, $p = 0.41$).




Figure 4.4 Decoding of landmark identity using MVPA. Landmark decoding accuracy (mean \pm SEM) within functionally and anatomically defined ROIs. Chance performance is 0.5. Hipp, Hippocampus. For accuracy by individual landmark, see Figure 4.7. *** $p < 0.001$.




Figure 4.5 Whole-brain (searchlight) analysis. Voxels in which landmark identity could be reliably decoded from response patterns in the surrounding neighborhood are plotted on an inflated version of the cortex. Light gray depicts gyri, and dark gray depicts sulci. Prototypical ROIs are overlaid for RSC (blue), PPA (green), and LOC (pink). These outlines were created by determining the average size of each ROI across subjects and plotting the across-subject ROI intersection that most closely matched that size. LH, left hemisphere; RH, right hemisphere.




Figure 4.6 Examples of the twenty-two images of one landmark, Huntsman Hall. Each image was shown only once in the experiment.

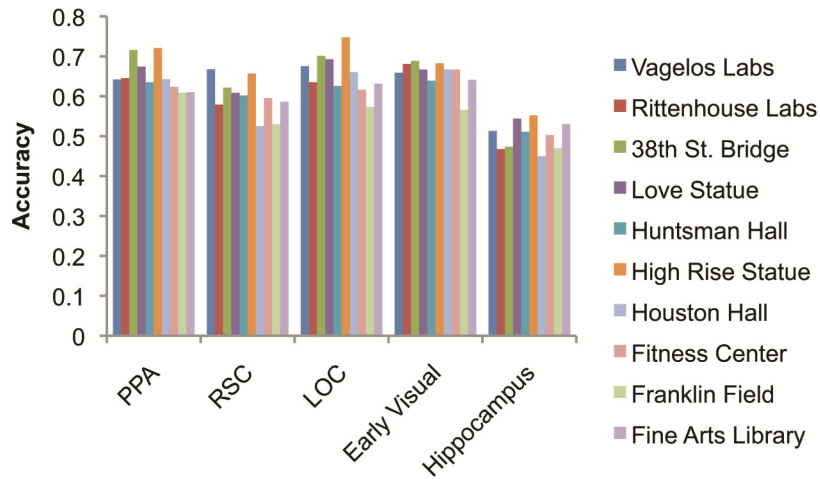


Figure 4.7 Decoding of individual landmarks using MVPA. Decoding accuracy within functionally and anatomically defined ROIs, displayed separately for each of the 10 landmarks. Chance performance is 0.5.

CHAPTER 5: GENERAL DISCUSSION AND CONCLUSIONS

Humans are able to navigate the world seemingly effortlessly, using information learned over multiple episodes of navigational experience to guide their current behavior. Although there has been extensive research on the neural basis of rodent navigation in small, simple environments, there has been a relative paucity of evidence for human spatial representations capable of supporting navigation through a large, complex environment like a city. To address this gap in the literature, I have conducted three fMRI experiments which aimed to investigate the neural representations of four types of spatial information that are critical to navigation: the navigator's current location (Chapters 2 & 3), the navigator's current allocentric heading (Chapters 2 & 3), the identities of landmarks (Chapter 4), and the distances between those landmarks (Chapter 4). To interrogate these representations, we scanned University of Pennsylvania students while they made decisions about places on the Penn campus. Consistent with predictions from the rodent neurophysiological data, and human neuropsychological and neuroimaging data, I observed spatial codes in regions of medial temporal and medial parietal cortex. I first describe how the main results from these three experiments fit with previous findings from neuropsychology, human behavior, computational modeling, neurophysiology, and neuroimaging. I then discuss how these results inform the functions of three specific regions of the human brain: parahippocampal place area (PPA), retrosplenial complex (RSC), and hippocampus. The experiments presented in Chapters 2 & 3 sought to identify regions of the human brain that represent the navigator's current location and heading in a large complex environment, a process known as self-localization. Prior neurophysiology work has largely implicated medial temporal lobe structures in this process (Jeffery, 2007; Barry and Burgess, 2014), but studies of human self-localization point to the importance of parahippocampal and retrosplenial regions (Epstein, 2008). We find support for the importance of all these regions in maintaining representations of location and heading. In Chapter 2, we measured activity patterns while subjects self-localized on the basis of

a photograph, and observed coding of location in RSC and left presubiculum and coding of heading in right presubiculum. In Chapter 3, we tested whether these spatial codes were consistent across perception and spatial imagery, and observed abstract coding of location in PPA and abstract coding of heading in RSC and entorhinal cortex (ERC). In Chapter 4, we interrogated representations of landmarks, and observed that PPA and RSC coded for landmark identity, as assessed by both fMRI adaptation (fMRIa) and multivoxel pattern analysis (MVPA). In contrast, left anterior hippocampus did not code for identity, but rather tracked the distance between landmarks in its bulk activity levels. Taken together, these findings have direct implications for a number of subfields in spatial cognition.

5.1 Relevance to neuropsychological literature

First, the current results provide insight into the neuropsychological literature and help explain the pattern of deficits observed in patients with focal brain lesions. The spatial codes we observed in PPA and RSC are broadly consistent with the behavioral profiles of patients with damage to parahippocampal and retrosplenial cortex respectively. The observation that patients with parahippocampal damage have difficulty recognizing prominent landmarks and environmental features (Pallis, 1955; Epstein et al., 2001; Takahashi and Kawamura, 2002; Mendez and Chierri, 2003) is consistent with our findings that PPA coded for landmark identity (Chapter 4) and for specific campus locations, especially when those locations were defined by either a single prominent landmark or a configuration of landmarks (Chapter 3). The observation that patients with retrosplenial damage cannot recall the directional relationships between landmarks is consistent with our findings that RSC coded for landmark identity (Chapter 4), specific campus locations (Chapter 2; N.B. a location can be uniquely defined by vectors to landmarks), and heading (Chapters 2 & 3), especially when headings were defined based on the directional relationship between landmarks (Chapter 3).

Furthermore, our results provide a potential explanation for one of the more puzzling findings in the neuropsychological literature: patients without functioning hippocampi retain the ability to navigate environments learned prior to injury (Teng and Squire, 1999; Rosenbaum et al., 2000; Maguire et al., 2006). Our results suggest that this residual environmental knowledge may be supported by regions such as PPA and RSC, which were capable of discriminating locations and headings at a fairly coarse resolution (i.e., locations separated by tens of meters, and headings separated by at least 90°). Indeed, the fact that we only observed fine-grained metric coding of location in the hippocampus (Chapter 4), and not PPA or RSC, is consistent with observations that hippocampal amnesics who lack this information have environmental representations that are less rich and detailed than those of control subjects (Rosenbaum et al., 2000; Maguire et al., 2006). Thus, our findings are consistent with the hypothesis that extra-hippocampal regions can support a coarse, schematic representation of space whereas the hippocampus is needed for a high fidelity, detailed spatial representation.

5.2 Relevance to human behavioral literature

Second, our results suggest a neural basis for the orientation specificity or “alignment effects” that have been observed in behavioral studies of spatial memory since at least the 1980s (Hintzman et al., 1981; see McNamara, 2003 for a review). In these studies, subjects are faster and/or more accurate when asked to retrieve spatial memories facing a particular direction. Although these experiments are most commonly performed using small tabletop environments, recent work has shown that large environmental spaces like the one studied here are also represented according to a particular reference direction (Marchette et al., 2011; Frankenstein et al., 2012). In Chapter 3, we replicated this behavioral effect and showed that subjects judged the direction of a target landmark significantly more quickly when they imagined or perceived a North-facing viewpoint. Critically, we observed that heading codes in RSC

exhibited the same pattern of orientation specificity; multivoxel patterns associated with North-facing views were significantly more distinguishable than views facing other directions. In contrast, the heading effects observed in ERC did not differ based on direction, suggesting that it does not contribute to the orientation-specific effects we observed.

One way to more conclusively demonstrate that RSC represents reference directions in spatial memory, would involve sampling heading representations at a finer scale, including headings that are misaligned with the principal axes of campus. This type of data would allow us to test whether heading representations in RSC follow the sawtooth pattern of results that has been observed for views facing non-preferred directions, a pattern characterized in the behavioral data by poorer performance for misaligned headings (e.g., NE, SE, SW, NW) than aligned headings (e.g., E, S, W; Marchette et al., 2011). Furthermore, since these alignment effects are also strong when subjects retrieve memories of newly learned environments (McNamara et al., 2003), RSC should code for the reference direction under this scenario as well. This is particularly useful because it means that these effects can be interrogated using virtual environments, which allow for full experimenter control of both visual and spatial properties.

5.3 Relevance to computational modeling

Third, our results confirm a prediction of a prominent model of spatial memory, which hypothesized that the regions involved in representing space during perception can also be driven by internally-generated mental imagery (Byrne et al., 2007). The results presented in Chapter 3 demonstrated consistent representations of location in PPA and consistent representations of heading in RSC and ERC when subjects were asked to retrieve directional information via bottom-up perceptual processes or top-down imagery processes. Furthermore, the model also appears to predict the finding that RSC only consistently codes the reference direction of the environment (North)

across perception and imagery. In the model, RSC serves as the transformation circuit to allow egocentric representations to be converted to allocentric representations and vice versa. When the circuit is driven by bottom-up inputs, independent populations of RSC neurons represent the egocentric information at a variety of orientations, each corresponding to the allocentric heading preferred by that population of neurons. However, when the circuit is driven by top-down inputs, all populations of RSC neurons maintain the same allocentric representation, which is then rotated to the appropriate egocentric view in precuneus. In other words, during bottom-up input, RSC is capable of representing all allocentric headings, but during top-down input, RSC only represents the heading that corresponds to the reference direction of the environment. As a result, when subjects performed both versions of the JRD task, only trials facing the reference direction would have elicited consistent representations across perception and imagery. Under this hypothesis, I would expect that if I measured the similarity of heading signals in RSC between trials of the same task, I would observe coding of all headings for picture (i.e., bottom-up) trials, but only coding of the reference direction for verbal (i.e., top-down) trials.

Although our RSC results seem to be consistent with the model's predictions, it is not entirely clear whether the same is true for our PPA and parahippocampal cortex (PHC) results. The model assumes that allocentric representations are supported by a population of boundary vector cells (BVCs) in PHC, which fire whenever a boundary is at a particular distance and allocentric direction from the navigator. Although we did not explicitly test for BVC-like representations, one would expect based on their firing properties that they should support coding of discrete locations, insofar as those locations can be distinguished by differences in the arrangement of local boundaries. Indeed, in the model, this is the mechanism by which place cells acquire their place fields — on the basis of inputs from a set of BVCs. However, there are a few pieces of evidence which taken together suggest that PPA is not the locus of the BVC representation. Although PPA, and to a lesser extent PHC, distinguished between

locations in Chapter 3, the effects in PHC appeared to be largely driven by coding of the particular view. PPA also exhibited strong view coding in Chapter 2, whereas location coding was relatively weak and no longer significant when low-level visual similarity was accounted for. Furthermore, PPA coded for landmark identity in Chapter 4; critically, different views of the same landmark contained overlapping visual information, but did not depict a consistent allocentric relationship between the observer and the landmark, which would be required to elicit a stable representation in a population of BVCs. Finally, this computational model was developed before the existence of BVCs had been confirmed experimentally. To date, they have only been observed in the subiculum, not in PHC (Lever et al., 2009). Taken together, these results suggest that PPA and PHC may support a visuospatial representation rather than the strictly spatial representation predicted by the model. In the model, this visual information is assumed to be encoded by perirhinal cortex (PRC), a region which did not exhibit spatial coding in any of our experiments, and which is typically implicated in coding of object information rather than spatial or contextual information (Ranganath and Ritchey, 2012). Thus, the functions ascribed to PRC in the model may be more in line with the representations we observed in PPA/PHC. These regions might also support a BVC-like representation of space, but this will need to be confirmed using environments for which the distances and directions to boundaries/landmarks can be carefully controlled (e.g., a virtual environment).

5.4 Relevance to neurophysiological studies

Fourth, our results highlight the need for more neurophysiological recordings from parahippocampal, retrosplenial, and parietal cortex, especially in primates and humans. Putative homologous regions have been identified in rodents, but the extent of functional overlap between rodents and humans is presently unclear. For example, postrhinal cortex is the putative homolog of PHC, but lesions to this region do not result in navigation deficits (Burwell et al., 2004). Compared to other MTL regions, there have

been relatively few recordings from this area in rodents, but the available evidence points to a role in processing spatial context and linking objects to a particular place in the environment (Ho and Burwell, 2014). Likewise, only a few studies have investigated the spatial correlates of neurons in rodent posterior parietal cortex. So far, the results indicate that neurons in this region are capable of representing space in multiple reference frames (Nitz, 2006, 2012), and that the representations may be more tied to the behavioral state of the animal (e.g., free foraging vs. stereotyped path-running) than the spatial structure of the environment (Whitlock et al., 2012). In sum, much more neurophysiology work is needed before one can make strong claims about functional homologies between human PPA and RSC and regions in the rodent brain.

Although there have been many recordings from primate parietal cortex, the vast majority of these have been focused on understanding sensorimotor transformations—how information acquired in the reference frame of the sensory receptor (e.g., retinotopically, in the case of vision), such as the location of an object, is transformed into a reference frame suitable for motor output, like reaching to grasp the object (see Andersen et al., 1997 for a review). To my knowledge, there has only been one recording of macaque medial parietal neurons during navigation (Sato et al., 2006, 2010). When monkeys followed routes in a virtual environment, neurons in this region exhibited a variety of navigation-associated responses and were differentially sensitive to the route, location, and direction of movement, broadly similar to what has been observed in rodent recordings and to the results presented here.

Recent studies indicate that we may soon know more about the firing properties of individual neurons in PPA and RSC. Two independent groups have identified scene-selective regions in the macaque brain that appear to be homologous to human PPA and RSC and code for both spatial and nonspatial aspects of scenes during perception (Nasr et al., 2011; Kornblith et al., 2013). However, these regions have not yet been studied during navigation or other spatial memory tasks so the specific spatial firing correlates of these neurons is presently unknown.

5.5 Relevance to neuroimaging studies

Fifth, our results are directly relevant to neuroimaging studies of human spatial representations. As described in the introductory chapter, previous work has suggested that parahippocampal, retrosplenial, and medial temporal regions are all involved in navigation, but few studies have examined the spatial representations supported by these regions. Multivoxel pattern analysis (MVPA) has been previously used to show that the hippocampus represents position within a small, simple virtual environment (Hassabis et al., 2009; Rodriguez, 2010); however the same technique could not decode position within a larger, more complex environment (Op de Beeck et al., 2013). We were also unable to use hippocampal patterns to decode location within a large, real-world environment (Chapters 2 & 3). Instead, we were able to decode location in left presubiculum (Chapter 2), RSC (Chapter 2), and PPA (Chapter 3).

However, the absence of significant decoding does not necessarily indicate that the hippocampus does not contain location information. Indeed, using fMRI adaptation (fMRIa), we identified metric coding of location in the left anterior hippocampus (Chapter 4), a feature of the representation that was not observable in hippocampal multivoxel patterns. Recently, that finding was replicated by an independent group, which showed that signals in the anterior hippocampus tracked the Euclidean distance to a goal location, whereas signals in the posterior hippocampus tracked the route distance to that goal (Howard et al., 2014). Thus, results from our lab and others (e.g., Baumann and Mattingley, 2013; Howard et al., 2014) indicate that hippocampal spatial representations might be best elucidated by indexing changes in the bulk BOLD activity rather than measuring multivoxel pattern similarity (but see Chadwick et al., 2012). In general, the results from Chapters 2 and 4, which simultaneously measured representations using MVPA and fMRIa, indicate that these two techniques often lead to different results and suggest that these analyses may interrogate different aspects of the neural code (for further discussion on this topic, see Epstein and Morgan, 2012).

Previous work using fMRIa to study spatial representations found that medial parietal cortex represents allocentric heading within a virtual maze (Baumann and Mattingley, 2010), and that ERC exhibits activity consistent with a population of heading-modulated grid cells (Doeller et al., 2010). The results from Chapters 2 and 3 are generally consistent with this prior literature. We observed coding of heading in RSC as assessed by fMRIa (Chapter 2; N.B. we observed anti-adaptation rather than a classic adaptation response) and by MVPA (Chapter 3). We also observed directional coding in ERC (Chapter 3), which could have been supported by a population of heading-modulated grid cells. Additionally, we observed multivoxel coding of heading in right presubiculum (Chapter 2), though these codes were not task-invariant, as shown in Chapter 3. This may indicate that neurons in this region are sensitive to remapping, a phenomenon which has been widely observed in neurophysiological recordings (Muller and Kubie, 1987; Leutgeb et al., 2005; Leutgeb and Leutgeb, 2014) and recently observed in a neuroimaging study of grid-related activity (Pape et al., 2011). Taken together, our results indicate that heading codes are present in both medial temporal and medial parietal cortex, and the results presented in Chapter 3 showed that these representations are distinct. Future studies will be needed to further characterize the differences in heading codes between these regions, including differential sensitivity to environmental features and task manipulations.

5.6 Representations in PPA, RSC, and Hippocampus

What information do PPA, RSC, and hippocampus represent in service of navigation? To answer this question, I will consider our results in the context of the broader spatial cognition literature and describe the hypothesized role of each region in turn.

As described in Chapter 1, prior neuropsychological data strongly implicates PPA in processing and recognizing visual scenes and landmarks. Furthermore, prior neuroimaging work suggests that PPA is sensitive to both spatial and nonspatial

aspects of scenes. Here, we have shown that PPA codes landmark identity (Chapter 4), specific campus views (Chapter 2), and discrete campus locations (weakly in Chapter 2, robustly in Chapter 3). Because we used photographs as stimuli for the experiments presented in Chapters 2 and 4, these results could be explained based on representations of visual or spatial information. The results from Chapter 3 show that PPA's representation is not completely explained by specific visual features in the stimulus, as there was no overlap in visual information between different JRDs from different tasks. However, this experiment could not rule out the possibility that the decoding we observed was due to similarity between the perceived views and mental imagery of those views. Indeed, a recent experiment has shown that PPA can cross-classify scenes across perception and imagery (Johnson and Johnson, 2014).

The experiments presented here were not explicitly designed to differentiate between the specific features of familiar scenes that give rise to consistent representations in PPA. However, one possibility supported by the literature is that there are two different kinds of representations supported by different PPA subregions. Functional connectivity analyses indicate that posterior PPA is more strongly connected with low-level visual regions whereas anterior PPA is more strongly connected with RSC and other regions of the default mode network (Baldassano et al., 2013). Thus, posterior PPA may be more involved in representing the visuospatial components of the scene, such as its spatial envelope (Oliva and Torralba, 2001) or its spatially organized textures (Oliva and Torralba, 2007), two scene features that PPA has been shown to be sensitive to (Kravitz et al., 2011a; Cant and Xu, 2012; Park et al., 2014). Anterior PPA might be more involved in linking this visuospatial scene representation to other nearby scenes (i.e., scene-scene association; N.B. "nearby" could potentially refer to semantic, visual, or spatial proximity) or objects (i.e., object-scene association), which would be consistent with theories implicating anterior parahippocampal cortex in coding contextual information (see Aminoff et al., 2013 for a review). Either of these accounts could potentially explain the PPA results we observed in Chapters 2-4.

The results we observed for RSC are consistent with theories that have implicated this region in situating the local scene within a broader context (Epstein, 2008; Ranganath and Ritchey, 2012) and supporting transformations between egocentric and allocentric reference frames (Byrne et al., 2007; Vann et al., 2009). We observed that RSC coded for landmark identity (Chapter 4), specific campus locations (Chapter 2), and allocentric headings (Chapters 2 & 3). All of these results, as well as the data from neuropsychological patients, could be explained by hypothesizing that RSC encodes the vectors between landmarks or vistas (Epstein and Vass, 2014).

There is evidence that the nature of these vectors might differ between dorsal and ventral medial parietal cortex. Dorsal regions in the precuneus (i.e., more dorsal than RSC) appear to be largely involved in representing transient, egocentric vectors, such as those calculated when tracking object locations during self-motion. For example, this region is more active when subjects are required to update object locations during path integration (Wolbers et al., 2008) or to imagine rotations of an object array (Jahn et al., 2012; Lambrey et al., 2012), and multivoxel patterns in this region code for the egocentric direction to target objects (Schindler and Bartels, 2013). Furthermore, this region is functionally connected via the angular gyrus to area V6 (Kravitz et al., 2011b), a visual region that appears to be involved in calculating both object-motion and self-motion on the basis of optic flow information (Pitzalis et al., 2012; Pitzalis et al., 2013).

In contrast, ventral regions of medial parietal cortex including RSC appear to be involved in tasks that require retrieving enduring, allocentric information or transforming between egocentric and allocentric reference frames. RSC is strongly activated when subjects make decisions based solely on allocentric knowledge, such as determining which of two landmarks is closer to a third landmark (Rosenbaum et al., 2004). It is also strongly engaged when subjects are asked to retrieve allocentric knowledge, such as heading or location, on the basis of egocentric visual information, i.e., an egocentric-allocentric transformation (Epstein et al., 2007). Likewise, RSC is strongly activated by tasks requiring an allocentric-egocentric transformation, such as imagined perspective

changes (Lambrey et al., 2012) and judgments of relative direction (Zhang et al., 2012). We show here that allocentric information is explicitly represented in RSC (Chapters 2 and 3) and that the reference direction in particular is coded in a way that is invariant to the processes used to access the representation (Chapter 3).

Our results indicate that allocentric information is also represented in the hippocampus, but that it takes a different form than in RSC. Whereas RSC appears to support a coarse, schematic representation that is capable of distinguishing between locations and headings, the hippocampus has access to metric information—the distances between familiar landmarks (Chapter 4). This metric information may be derived from grid cells, which provide input to the hippocampus and explicitly encode the distance between locations in nearby space. The fact that we detected metric information even though subjects were not asked to retrieve it (N.B. subjects were performing a simple recognition task) suggests that it may be automatically activated when thinking about familiar places. This is consistent with recent human intracranial recordings, which showed that when subjects are asked to recall the identities of objects encountered during a navigational epoch, they automatically retrieved the spatial locations associated with those objects (Miller et al., 2013). In sum, our results show that the hippocampus has access to detailed, metric spatial information, a key feature of cognitive maps.

How do PPA, RSC, and hippocampus work together to support spatial memory? I briefly describe one possible model that can tie the functions of these regions together (Figure 5.1). Precise spatial locations are coded in the hippocampus by a population of place cells. When long-term spatial memory is accessed, the place cells activate representations in PPA/PHC, which encode the associated visuospatial information, such as views visible from that location. This information is sent to RSC, which represents vectors to landmarks according to the reference direction of the environment. This may include vectors to landmarks visible from that location, as well as vectors to salient unseen landmarks, where saliency may be defined by visual,

semantic, or structural properties, such as being located at a prominent intersection (Richter and Winter, 2014). Online transformations of these vectors (i.e., translations and rotations) are computed in dorsal medial parietal regions such as precuneus, where the information can ultimately be converted into a reference frame appropriate for motor output.

5.7 Conclusions

The experiments presented in this dissertation sought to uncover the regions of the human brain that support long-term memory representations of large-scale environmental space. In designing these experiments and formulating our hypotheses, we drew inspiration from many different subfields of spatial cognition. Our results are largely consistent with this vast body of knowledge and serve to further elucidate the roles of specific regions of medial temporal and medial parietal cortex. Ultimately, this knowledge may serve as a framework for future experiments aimed at understanding how representations in these regions interact during spatial memory retrieval, are formed over the course of navigational experience, and might differ across individuals as a function of navigational aptitude.

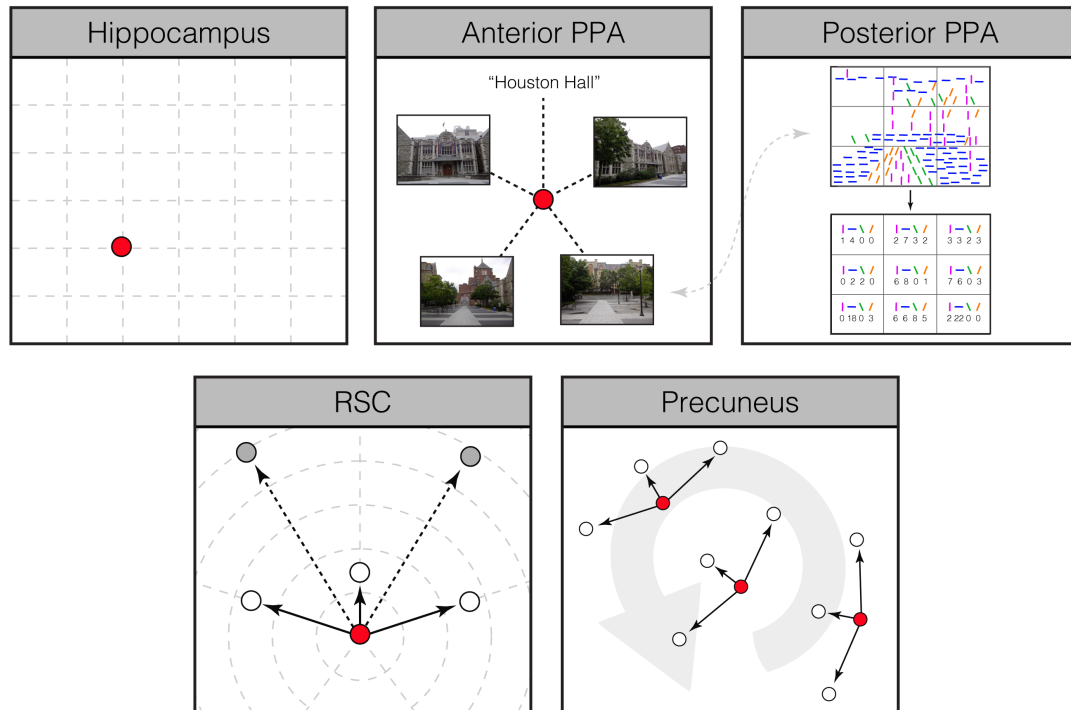


Figure 5.1 Model of medial temporal and medial parietal representations. The hippocampus contains a precise allocentric representation of the locations of landmarks in the environment (1 landmark shown in red). Anterior PPA contains links to the associated visuospatial and possibly semantic information. This could include the name of the landmark, different views of that landmark (top 2 photographs), or scenes that can be viewed from that landmark (bottom 2 photographs). Posterior PPA represents these views according to their visual and spatial features. One possible representation is a configuration of spatially organized textures, such as lines of different orientations (shown in different colors). This information might be organized at a relatively fine scale, as shown at top, or at a more coarse scale, as shown at bottom, where each region of space is described by summary statistics. For alternative scene representations, see Oliva and Torralba, 2007. RSC represents the vectors to landmarks within a coarse, polar map aligned to the reference direction of the environment. Vectors might indicate the locations of visible landmarks (white circles) or salient unseen landmarks (gray circles). Precuneus performs online rotations of the landmark vectors represented in RSC. To minimize the load, precuneus might only rotate vectors that are necessary to produce the appropriate motor behavior. For example, in this case, it does not rotate the vectors to unseen landmarks.

BIBLIOGRAPHY

- Aguirre GK (2007) Continuous carry-over designs for fMRI. *Neuroimage* 35:1480-1494.
- Aguirre GK, D'Esposito M (1997) Environmental knowledge is subserved by separable dorsal/ventral neural areas. *J Neurosci* 17:2512-2518.
- Aguirre GK, D'Esposito M (1999) Topographical disorientation: a synthesis and taxonomy. *Brain : a journal of neurology* 122 (Pt 9):1613-1628.
- Aguirre GK, Detre JA, Alsop DC, D'Esposito M (1996) The parahippocampus subserves topographical learning in man. *Cereb Cortex* 6:823-829.
- Alsaadi T, Binder JR, Lazar RM, Doorani T, Mohr JP (2000) Pure topographic disorientation: A distinctive syndrome with varied localization. *Neurology* 54:1864-1866.
- Aminoff EM, Kveraga K, Bar M (2013) The role of the parahippocampal cortex in cognition. *Trends Cogn Sci* 17:379-390.
- Andersen RA, Snyder LH, Bradley DC, Xing J (1997) Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annual review of neuroscience* 20:303-330.
- Baldassano C, Beck DM, Fei-Fei L (2013) Differential connectivity within the Parahippocampal Place Area. *Neuroimage* 75:228-237.
- Barry C, Burgess N (2014) Neural mechanisms of self-location. *Current biology : CB* 24:R330-339.
- Baumann O, Mattingley JB (2010) Medial parietal cortex encodes perceived heading direction in humans. *J Neurosci* 30:12897-12901.
- Baumann O, Mattingley JB (2013) Dissociable representations of environmental size and complexity in the human hippocampus. *J Neurosci* 33:10526-10533.
- Boccaro CN, Sargolini F, Thoresen VH, Solstad T, Witter MP, Moser EI, Moser MB (2010) Grid cells in pre- and parasubiculum. *Nat Neurosci* 13:987-994.
- Brown MW, Wilson FA, Riches IP (1987) Neuronal evidence that inferomedial temporal cortex is more important than hippocampus in certain processes underlying recognition memory. *Brain research* 409:158-162.
- Brown TI, Ross RS, Keller JB, Hasselmo ME, Stern CE (2010) Which way was I going? Contextual retrieval supports the disambiguation of well learned overlapping navigational routes. *J Neurosci* 30:7414-7422.
- Burgess N, Jackson A, Hartley T, O'Keefe J (2000) Predictions derived from modelling the hippocampal role in navigation. *Biological cybernetics* 83:301-312.
- Burwell RD, Saddoris MP, Bucci DJ, Wiig KA (2004) Corticohippocampal contributions to spatial and contextual learning. *J Neurosci* 24:3826-3836.
- Byrne P, Becker S, Burgess N (2007) Remembering the past and imagining the future: a neural model of spatial memory and imagery. *Psychological review* 114:340-375.
- Cacucci F, Lever C, Wills TJ, Burgess N, O'Keefe J (2004) Theta-modulated place-by-direction cells in the hippocampal formation in the rat. *J Neurosci* 24:8265-8277.
- Caglio M, Castelli L, Cerrato P, Latini-Corazzini L (2011) Pure topographical disorientation in a patient with right occipito-temporal lesion. *Neurosciences* 16:60-65.
- Cant JS, Xu Y (2012) Object ensemble processing in human anterior-medial ventral visual cortex. *J Neurosci* 32:7685-7700.
- Chadwick MJ, Bonnici HM, Maguire EA (2012) Decoding information in the human hippocampus: a user's guide. *Neuropsychologia* 50:3107-3121.
- Chadwick MJ, Mullally SL, Maguire EA (2013) The hippocampus extrapolates beyond the view in scenes: an fMRI study of boundary extension. *Cortex* 49:2067-2079.
- Chadwick MJ, Hassabis D, Weiskopf N, Maguire EA (2010) Decoding individual episodic memory traces in the human hippocampus. *Current biology : CB* 20:544-547.

- Chen LL, Lin LH, Green EJ, Barnes CA, McNaughton BL (1994) Head-direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. *Experimental brain research* 101:8-23.
- Cho J, Sharp PE (2001) Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behavioral neuroscience* 115:3-25.
- Committeri G, Galati G, Paradis AL, Pizzamiglio L, Berthoz A, LeBihan D (2004) Reference frames for spatial cognition: different brain areas are involved in viewer-, object-, and landmark-centered judgments about object location. *J Cogn Neurosci* 16:1517-1535.
- Corkin S (2002) What's new with the amnesic patient H.M.? *Nature reviews Neuroscience* 3:153-160.
- Deshmukh SS, Knierim JJ (2013) Influence of local objects on hippocampal representations: Landmark vectors and memory. *Hippocampus* 23:253-267.
- DiCarlo JJ, Cox DD (2007) Untangling invariant object recognition. *Trends Cogn Sci* 11:333-341.
- Dilks DD, Julian JB, Kubilius J, Spelke ES, Kanwisher N (2011) Mirror-image sensitivity and invariance in object and scene processing pathways. *J Neurosci* 31:11305-11312.
- Doeller CF, Barry C, Burgess N (2010) Evidence for grid cells in a human memory network. *Nature* 463:657-661.
- Drucker DM, Aguirre GK (2009) Different spatial scales of shape similarity representation in lateral and ventral LOC. *Cereb Cortex* 19:2269-2280.
- Ekstrom AD, Kahana MJ, Caplan JB, Fields TA, Isham EA, Newman EL, Fried I (2003) Cellular networks underlying human spatial navigation. *Nature* 425:184-188.
- Epstein R, Kanwisher N (1998) A cortical representation of the local visual environment. *Nature* 392:598-601.
- Epstein R, Graham KS, Downing PE (2003) Viewpoint-specific scene representations in human parahippocampal cortex. *Neuron* 37:865-876.
- Epstein R, Deyoe EA, Press DZ, Rosen AC, Kanwisher N (2001) Neuropsychological evidence for a topographical learning mechanism in parahippocampal cortex. *Cognitive neuropsychology* 18:481-508.
- Epstein RA (2008) Parahippocampal and retrosplenial contributions to human spatial navigation. *Trends Cogn Sci* 12:388-396.
- Epstein RA, Morgan LK (2012) Neural responses to visual scenes reveals inconsistencies between fMRI adaptation and multivoxel pattern analysis. *Neuropsychologia* 50:530-543.
- Epstein RA, Vass LK (2014) Neural systems for landmark-based wayfinding in humans. *Philosophical transactions of the Royal Society of London Series B, Biological sciences* 369:20120533.
- Epstein RA, Parker WE, Feiler AM (2007) Where am I now? Distinct roles for parahippocampal and retrosplenial cortices in place recognition. *J Neurosci* 27:6141-6149.
- Epstein RA, Parker WE, Feiler AM (2008) Two kinds of FMRI repetition suppression? Evidence for dissociable neural mechanisms. *Journal of neurophysiology* 99:2877-2886.
- Foster DJ, Wilson MA (2006) Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* 440:680-683.
- Frankenstein J, Mohler BJ, Bulthoff HH, Meilinger T (2012) Is the map in our head oriented north? *Psychological science* 23:120-125.
- Froehner MT, Duffy CJ (2002) Cortical neurons encoding path and place: where you go is where you are. *Science* 295:2462-2465.
- Fyhn M, Hafting T, Treves A, Moser MB, Moser EI (2007) Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* 446:190-194.
- Galati G, Pelle G, Berthoz A, Committeri G (2010) Multiple reference frames used by the human brain for spatial perception and memory. *Experimental brain research* 206:109-120.

- Gallistel CR (1990) *The Organization of Learning*. Cambridge, MA: MIT Press.
- Georges-Francois P, Rolls ET, Robertson RG (1999) Spatial view cells in the primate hippocampus: allocentric view not head direction or eye position or place. *Cereb Cortex* 9:197-212.
- Ghaem O, Mellet E, Crivello F, Tzourio N, Mazoyer B, Berthoz A, Denis M (1997) Mental navigation along memorized routes activates the hippocampus, precuneus, and insula. *Neuroreport* 8:739-744.
- Giocomo LM, Stensola T, Bonnevie T, Van Cauter T, Moser MB, Moser EI (2014) Topography of head direction cells in medial entorhinal cortex. *Current biology : CB* 24:252-262.
- Gray JA, McNaughton N (1982) *The neuropsychology of anxiety: an enquiry into the functions of the septo-hippocampal system*. New York: Oxford University Press.
- Grill-Spector K, Malach R (2001) fMR-adaptation: a tool for studying the functional properties of human cortical neurons. *Acta psychologica* 107:293-321.
- Grill-Spector K, Henson R, Martin A (2006) Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn Sci* 10:14-23.
- Habib M, Sirigu A (1987) Pure topographical disorientation: a definition and anatomical basis. *Cortex* 23:73-85.
- Hafting T, Fyhn M, Molden S, Moser MB, Moser EI (2005) Microstructure of a spatial map in the entorhinal cortex. *Nature* 436:801-806.
- Hartley T, Maguire EA, Spiers HJ, Burgess N (2003) The well-worn route and the path less traveled: distinct neural bases of route following and wayfinding in humans. *Neuron* 37:877-888.
- Hashimoto R, Tanaka Y, Nakano I (2010) Heading disorientation: a new test and a possible underlying mechanism. *Eur Neurol* 63:87-93.
- Hassabis D, Chu C, Rees G, Weiskopf N, Molyneux PD, Maguire EA (2009) Decoding neuronal ensembles in the human hippocampus. *Current biology : CB* 19:546-554.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293:2425-2430.
- Hintzman DL, O'Dell CS, Arndt DR (1981) Orientation in cognitive maps. *Cognitive psychology* 13:149-206.
- Ho JW, Burwell RD (2014) Perirhinal and postrhinal functional inputs to the hippocampus. In: *Space, Time and Memory in the Hippocampal Formation* (Derdikman D, Knierim JJ, eds), pp 55-81. New York: Springer.
- Howard LR, Javadi AH, Yu Y, Mill RD, Morrison LC, Knight R, Loftus MM, Staskute L, Spiers HJ (2014) The Hippocampus and Entorhinal Cortex Encode the Path and Euclidean Distances to Goals during Navigation. *Current biology : CB* 24:1331-1340.
- Hupbach A, Hardt O, Gomez R, Nadel L (2008) The dynamics of memory: context-dependent updating. *Learning & memory* 15:574-579.
- Jacobs J (2014) Hippocampal theta oscillations are slower in humans than in rodents: implications for models of spatial navigation and memory. *Philosophical transactions of the Royal Society of London Series B, Biological sciences* 369:20130304.
- Jacobs J, Kahana MJ (2010) Direct brain recordings fuel advances in cognitive electrophysiology. *Trends Cogn Sci* 14:162-171.
- Jacobs J, Kahana MJ, Ekstrom AD, Mollison MV, Fried I (2010) A sense of direction in human entorhinal cortex. *Proc Natl Acad Sci U S A* 107:6487-6492.
- Jacobs J, Weidemann CT, Miller JF, Solway A, Burke JF, Wei XX, Suthana N, Sperling MR, Sharan AD, Fried I, Kahana MJ (2013) Direct recordings of grid-like neuronal activity in human spatial navigation. *Nat Neurosci* 16:1188-1190.
- Jahn G, Wendt J, Lotze M, Papenmeier F, Huff M (2012) Brain activation during spatial updating and attentive tracking of moving targets. *Brain and cognition* 78:105-113.

- Jeffery KJ (2007) Self-localization and the entorhinal-hippocampal system. *Current opinion in neurobiology* 17:684-691.
- Jeffery KJ (2013) Neural odometry: the discrete charm of the entorhinal cortex. *Current biology* : CB 23:R204-206.
- Jeffery KJ, Burgess N (2006) A metric for the cognitive map: found at last? *Trends Cogn Sci* 10:1-3.
- Jenkinson M, Smith S (2001) A global optimisation method for robust affine registration of brain images. *Medical image analysis* 5:143-156.
- Jenkinson M, Bannister P, Brady M, Smith S (2002) Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* 17:825-841.
- Jenkinson M, Beckmann CF, Behrens TE, Woolrich MW, Smith SM (2012) Fsl. *Neuroimage* 62:782-790.
- Johnson A, Redish AD (2007) Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci* 27:12176-12189.
- Johnson MR, Johnson MK (2014) Decoding individual natural scene representations during perception and imagery. *Frontiers in human neuroscience* 8:59.
- Julian JB, Fedorenko E, Webster J, Kanwisher N (2012) An algorithmic method for functionally defining regions of interest in the ventral visual pathway. *Neuroimage* 60:2357-2364.
- Kasai K, Shenton ME, Salisbury DF, Onitsuka T, Toner SK, Yurgelun-Todd D, Kikinis R, Jolesz FA, McCarley RW (2003) Differences and similarities in insular and temporal pole MRI gray matter volume abnormalities in first-episode schizophrenia and affective psychosis. *Archives of general psychiatry* 60:1069-1077.
- Katayama K, Takahashi N, Ogawara K, Hattori T (1999) Pure Topographical Disorientation Due to Right Posterior Cingulate Lesion. *Cortex* 35:279-282.
- Kim JJ, Crespo-Facorro B, Andreasen NC, O'Leary DS, Zhang B, Harris G, Magnotta VA (2000) An MRI-based parcellation method for the temporal lobe. *Neuroimage* 11:271-288.
- Kjelstrup KB, Solstad T, Brun VH, Hafting T, Leutgeb S, Witter MP, Moser EI, Moser MB (2008) Finite scale of spatial representation in the hippocampus. *Science* 321:140-143.
- Klatzky RL (1998) Allocentric and egocentric spatial representations: definitions, distinctions, and interconnections. In: *Spatial Cognition: An Interdisciplinary Approach to Representing and Processing Spatial Knowledge* (Freksa C, Habel C, Wender KF, eds). New York: Springer.
- Knierim JJ (2002) Dynamic interactions between local surface cues, distal landmarks, and intrinsic circuitry in hippocampal place cells. *J Neurosci* 22:6254-6264.
- Knierim JJ, Hamilton DA (2011) Framing spatial cognition: neural representations of proximal and distal frames of reference and their roles in navigation. *Physiological reviews* 91:1245-1279.
- Kornblith S, Cheng X, Ohayon S, Tsao DY (2013) A network for scene processing in the macaque temporal lobe. *Neuron* 79:766-781.
- Kourtzi Z, Kanwisher N (2001) Representation of perceived object shape by the human lateral occipital complex. *Science* 293:1506-1509.
- Kourtzi Z, Bulthoff HH, Erb M, Grodd W (2002) Object-selective responses in the human motion area MT/MST. *Nat Neurosci* 5:17-18.
- Kravitz DJ, Peng CS, Baker CI (2011a) Real-world scene representations in high-level visual cortex: it's the spaces more than the places. *J Neurosci* 31:7322-7333.
- Kravitz DJ, Saleem KS, Baker CI, Mishkin M (2011b) A new neural framework for visuospatial processing. *Nature reviews Neuroscience* 12:217-230.
- Kriegeskorte N, Goebel R, Bandettini P (2006) Information-based functional brain mapping. *Proc Natl Acad Sci U S A* 103:3863-3868.
- Kubie JL, Fenton AA (2012) Linear look-ahead in conjunctive cells: an entorhinal mechanism for vector-based navigation. *Frontiers in neural circuits* 6:20.

- Kumaran D, Maguire EA (2006) An unexpected sequence of events: mismatch detection in the human hippocampus. *PLoS biology* 4:e424.
- Kumaran D, Maguire EA (2007) Which computational mechanisms operate in the hippocampus during novelty detection? *Hippocampus* 17:735-748.
- Lambrey S, Doeller C, Berthoz A, Burgess N (2012) Imagining being somewhere else: neural basis of changing perspective in space. *Cereb Cortex* 22:166-174.
- Larsson J, Smith AT (2012) fMRI repetition suppression: neuronal adaptation or stimulus expectation? *Cereb Cortex* 22:567-576.
- Leutgeb S, Leutgeb JK (2014) Remapping to discriminate contexts with hippocampal population codes. In: *Space, Time, and Memory in the Hippocampal Formation* (Derdikman D, Knierim JJ, eds). New York: Springer.
- Leutgeb S, Leutgeb JK, Barnes CA, Moser EI, McNaughton BL, Moser MB (2005) Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* 309:619-623.
- Lever C, Burton S, Jeewajee A, O'Keefe J, Burgess N (2009) Boundary vector cells in the subiculum of the hippocampal formation. *J Neurosci* 29:9771-9777.
- Luzzi S, Pucci E, Di Bella P, Piccirilli M (2000) Topographical Disorientation Consequent to Amnesia of Spatial Location in A Patient with Right Parahippocampal Damage. *Cortex* 36:427-434.
- Maguire EA (2001) The retrosplenial contribution to human navigation: a review of lesion and neuroimaging findings. *Scand J Psychol* 42:225-238.
- Maguire EA, Frackowiak RS, Frith CD (1997) Recalling routes around london: activation of the right hippocampus in taxi drivers. *J Neurosci* 17:7103-7110.
- Maguire EA, Nannery R, Spiers HJ (2006) Navigation around London by a taxi driver with bilateral hippocampal lesions. *Brain : a journal of neurology* 129:2894-2907.
- Maguire EA, Burgess N, Donnett JG, Frackowiak RS, Frith CD, O'Keefe J (1998) Knowing where and getting there: a human navigation network. *Science* 280:921-924.
- Manns JR, Eichenbaum H (2009) A cognitive map for object memory in the hippocampus. *Learning & memory* 16:616-624.
- Marchette SA, Yerramsetti A, Burns TJ, Shelton AL (2011) Spatial memory in the real world: long-term representations of everyday environments. *Memory & cognition* 39:1401-1408.
- Marchette SA, Vass LK, Ryan J, Epstein RA (2014) For familiar landmarks, parahippocampal cortex represents place identity, not just perceptual features. In: *Vision Sciences Society*. St. Pete Beach, FL.
- Matsumoto H, Simmons A, Williams S, Hadjulis M, Pipe R, Murray R, Frangou S (2001) Superior temporal gyrus abnormalities in early-onset schizophrenia: similarities and differences with adult-onset schizophrenia. *The American journal of psychiatry* 158:1299-1304.
- Matsumura N, Nishijo H, Tamura R, Eifuku S, Endo S, Ono T (1999) Spatial- and task-dependent neuronal responses during real and virtual translocation in the monkey hippocampal formation. *J Neurosci* 19:2381-2393.
- McNamara TP (2003) How are the locations of objects in the environment represented in memory? In: *Spatial cognition III: Routes and navigation, human memory and learning, spatial representation and spatial reasoning* (Freksa C, Brauer, W., Habel, C., Wender, K., ed), pp 174-191. Berlin: Springer-Verlag.
- McNamara TP, Rump B, Werner S (2003) Egocentric and geocentric frames of reference in memory of large-scale space. *Psychon Bull Rev* 10:589-595.
- Mendez MF, Cherrier MM (2003) Agnosia for scenes in topographagnosia. *Neuropsychologia* 41:1387-1395.
- Miller JF, Neufang M, Solway A, Brandt A, Trippel M, Mader I, Hefft S, Merkow M, Polyn SM, Jacobs J, Kahana MJ, Schulze-Bonhage A (2013) Neural activity in human hippocampal formation reveals the spatial context of retrieved memories. *Science* 342:1111-1114.

- Montello DR, Hegarty M, Richardson AE, Waller D (2004) Spatial memory of real environments, virtual environments, and maps. In: Human spatial memory: Remembering where (Allen GL, ed), pp 251-285. Mahwah, NJ: Lawrence Erlbaum Associates.
- Morgan LK, Macevoy SP, Aguirre GK, Epstein RA (2011) Distances between real-world locations are represented in the human hippocampus. *J Neurosci* 31:1238-1245.
- Morris RG, Garrud P, Rawlins JN, O'Keefe J (1982) Place navigation impaired in rats with hippocampal lesions. *Nature* 297:681-683.
- Mou W, McNamara TP, Valiquette CM, Rump B (2004) Allocentric and egocentric updating of spatial memories. *Journal of experimental psychology Learning, memory, and cognition* 30:142-157.
- Muller RU, Kubie JL (1987) The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J Neurosci* 7:1951-1968.
- Mur M, Bandettini PA, Kriegeskorte N (2009) Revealing representational content with pattern-information fMRI--an introductory guide. *Social cognitive and affective neuroscience* 4:101-109.
- Nadel L, Willner J (1980) Context and conditioning: a place for space. *Physiol Psychol* 8:218-228.
- Nasr S, Liu N, Devaney KJ, Yue X, Rajimehr R, Ungerleider LG, Tootell RB (2011) Scene-selective cortical regions in human and nonhuman primates. *J Neurosci* 31:13771-13785.
- Navratilova Z, McNaughton BL (2014) Models of path integration in the hippocampal complex. In: *Space, Time, and Memory in the Hippocampal Formation* (Derdikman D, Knierim JJ, eds), pp 191-218. New York: Springer.
- Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Human brain mapping* 15:1-25.
- Nili H, Wingfield C, Walther A, Su L, Marslen-Wilson W, Kriegeskorte N (2014) A toolbox for representational similarity analysis. *PLoS computational biology* 10:e1003553.
- Nitz DA (2006) Tracking route progression in the posterior parietal cortex. *Neuron* 49:747-756.
- Nitz DA (2012) Spaces within spaces: rat parietal cortex neurons register position across three reference frames. *Nat Neurosci* 15:1365-1367.
- Norman KA, Polyn SM, Detre GJ, Haxby JV (2006) Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10:424-430.
- O'Keefe J (1976) Place units in the hippocampus of the freely moving rat. *Experimental neurology* 51:78-109.
- O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain research* 34:171-175.
- O'Keefe J, Conway DH (1978) Hippocampal place units in the freely moving rat: why they fire where they fire. *Experimental brain research* 31:573-590.
- O'Keefe J, Nadel L (1978) *The Hippocampus as a Cognitive Map*. Oxford: Oxford University Press.
- O'Keefe J, Burgess N (1996) Geometric determinants of the place fields of hippocampal neurons. *Nature* 381:425-428.
- Oliva A, Torralba A (2001) Modeling the shape of the scene: a holistic representation of the spatial envelope. *International Journal of Computer Vision* 42:145-175.
- Oliva A, Torralba A (2007) The role of context in object recognition. *Trends Cogn Sci* 11:520-527.
- Op de Beeck HP, Vermaercke B, Woolley DG, Wenderoth N (2013) Combinatorial brain decoding of people's whereabouts during visuospatial navigation. *Frontiers in neuroscience* 7:78.
- Osawa A, Maeshima S, Kunishio K (2008) Topographic disorientation and amnesia due to cerebral hemorrhage in the left retrosplenial region. *Eur Neurol* 59:79-82.

- Pallis CA (1955) Impaired identification of faces and places with agnosia for colours; report of a case due to cerebral embolism. *Journal of neurology, neurosurgery, and psychiatry* 18:218-224.
- Pape A-A, Wolbers T, Schultz J, Bulthoff HH, Meilinger T (2011) Grid cell remapping in humans. In: *Society for Neuroscience*. Washington, D.C.
- Park S, Chun MM (2009) Different roles of the parahippocampal place area (PPA) and retrosplenial cortex (RSC) in panoramic scene perception. *Neuroimage* 47:1747-1756.
- Park S, Konkle T, Oliva A (2014) Parametric Coding of the Size and Clutter of Natural Scenes in the Human Brain. *Cereb Cortex*.
- Park S, Brady TF, Greene MR, Oliva A (2011) Disentangling scene content from spatial boundary: complementary roles for the parahippocampal place area and lateral occipital complex in representing real-world scenes. *J Neurosci* 31:1333-1340.
- Park S, Intraub H, Yi DJ, Widders D, Chun MM (2007) Beyond the edges of a view: boundary extension in human scene-selective visual cortex. *Neuron* 54:335-342.
- Pitzalis S, Fattori P, Galletti C (2012) The functional role of the medial motion area V6. *Frontiers in behavioral neuroscience* 6:91.
- Pitzalis S, Sdoia S, Bultrini A, Committeri G, Di Russo F, Fattori P, Galletti C, Galati G (2013) Selectivity to translational egomotion in human brain motion areas. *PloS one* 8:e60241.
- Pruessner JC, Kohler S, Crane J, Pruessner M, Lord C, Byrne A, Kabani N, Collins DL, Evans AC (2002) Volumetry of temporopolar, perirhinal, entorhinal and parahippocampal cortex from high-resolution MR images: considering the variability of the collateral sulcus. *Cereb Cortex* 12:1342-1353.
- Quirk GJ, Muller RU, Kubie JL (1990) The firing of hippocampal place cells in the dark depends on the rat's recent experience. *J Neurosci* 10:2008-2017.
- Ranganath C, Ritchey M (2012) Two cortical systems for memory-guided behaviour. *Nature reviews Neuroscience* 13:713-726.
- Redish AD, Battaglia FP, Chawla MK, Ekstrom AD, Gerrard JL, Lipa P, Rosenzweig ES, Worley PF, Guzowski JF, McNaughton BL, Barnes CA (2001) Independence of firing correlates of anatomically proximate hippocampal pyramidal cells. *J Neurosci* 21:RC134.
- Renninger LW, Malik J (2004) When is scene identification just texture recognition? *Vision Res* 44:2301-2311.
- Richter K-F, Winter S (2014) Computational aspects: how landmarks can be observed, stored, and analysed. In: *Landmarks*, pp 137-173. New York: Springer.
- Robertson RG, Rolls ET, Georges-Francois P, Panzeri S (1999) Head direction cells in the primate pre-subiculum. *Hippocampus* 9:206-219.
- Rodriguez PF (2010) Neural decoding of goal locations in spatial navigation in humans with fMRI. *Human brain mapping* 31:391-397.
- Rolls ET (1999) Spatial view cells and the representation of place in the primate hippocampus. *Hippocampus* 9:467-480.
- Rosenbaum RS, Ziegler M, Winocur G, Grady CL, Moscovitch M (2004) "I have often walked down this street before": fMRI studies on the hippocampus and other structures during mental navigation of an old environment. *Hippocampus* 14:826-835.
- Rosenbaum RS, Priselac S, Kohler S, Black SE, Gao F, Nadel L, Moscovitch M (2000) Remote spatial memory in an amnesic person with extensive bilateral hippocampal lesions. *Nat Neurosci* 3:1044-1048.
- Sapountzis P, Schluppeck D, Bowtell R, Peirce JW (2010) A comparison of fMRI adaptation and multivariate pattern classification analysis in visual cortex. *Neuroimage* 49:1632-1640.
- Sargolini F, Fyhn M, Hafting T, McNaughton BL, Witter MP, Moser MB, Moser EI (2006) Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* 312:758-762.

- Sato N, Sakata H, Tanaka YL, Taira M (2006) Navigation-associated medial parietal neurons in monkeys. *Proc Natl Acad Sci U S A* 103:17001-17006.
- Sato N, Sakata H, Tanaka YL, Taira M (2010) Context-dependent place-selective responses of the neurons in the medial parietal region of macaque monkeys. *Cereb Cortex* 20:846-858.
- Sawamura H, Orban GA, Vogels R (2006) Selectivity of neuronal adaptation does not match response selectivity: a single-cell study of the fMRI adaptation paradigm. *Neuron* 49:307-318.
- Schinazi VR, Epstein RA (2010) Neural correlates of real-world route learning. *Neuroimage* 53:725-735.
- Schindler A, Bartels A (2013) Parietal cortex codes for egocentric space beyond the field of view. *Current biology* : CB 23:177-182.
- Segaert K, Weber K, de Lange FP, Petersson KM, Hagoort P (2013) The suppression of repetition enhancement: A review of fMRI studies. *Neuropsychologia* 51:59-66.
- Shapiro ML, Tanila H, Eichenbaum H (1997) Cues that hippocampal place cells encode: dynamic and hierarchical representation of local and distal stimuli. *Hippocampus* 7:624-642.
- Shelton AL, McNamara TP (1997) Multiple views of spatial memory. *Psychonomic Bulletin & Review* 4:102-106.
- Shelton AL, Gabrieli JD (2002) Neural correlates of encoding space from route and survey perspectives. *J Neurosci* 22:2711-2717.
- Sholl MJ (2001) The role of a self-reference system in spatial navigation. In: *Spatial Information Theory: Foundations of Geographic Information Science International Conference, COSIT 2001 Morrow Bay, CA, USA, September 19-23, 2001 Proceedings* (Montello DR, ed). New York: Springer.
- Shrager Y, Kirwan CB, Squire LR (2008) Neural basis of the cognitive map: path integration does not require hippocampus or entorhinal cortex. *Proc Natl Acad Sci U S A* 105:12034-12038.
- Smith DM, Barredo J, Mizumori SJ (2012) Complimentary roles of the hippocampus and retrosplenial cortex in behavioral context discrimination. *Hippocampus* 22:1121-1133.
- Solstad T, Boccara CN, Kropff E, Moser MB, Moser EI (2008) Representation of geometric borders in the entorhinal cortex. *Science* 322:1865-1868.
- Spiers HJ, Maguire EA (2006) Thoughts, behaviour, and brain dynamics during navigation in the real world. *Neuroimage* 31:1826-1840.
- Spiers HJ, Maguire EA (2007) A navigational guidance system in the human brain. *Hippocampus* 17:618-626.
- Spreng RN, Mar RA, Kim AS (2009) The common neural basis of autobiographical memory, prospection, navigation, theory of mind, and the default mode: a quantitative meta-analysis. *J Cogn Neurosci* 21:489-510.
- Squire LR (1992) Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychological review* 99:195-231.
- Stensola H, Stensola T, Solstad T, Froland K, Moser MB, Moser EI (2012) The entorhinal grid map is discretized. *Nature* 492:72-78.
- Sternberg S (1998) Inferring mental operations from reaction-time data: How we compare objects. In: *Methods, models, and conceptual issues*, 2 Edition (Scarborough D, Sternberg S, eds). Cambridge, MA: MIT Press.
- Strange BA, Dolan RJ (2001) Adaptive anterior hippocampal responses to oddball stimuli. *Hippocampus* 11:690-698.
- Summerfield C, Trittschuh EH, Monti JM, Mesulam MM, Egner T (2008) Neural repetition suppression reflects fulfilled perceptual expectations. *Nat Neurosci* 11:1004-1006.
- Suthana NA, Ekstrom AD, Moshirvaziri S, Knowlton B, Bookheimer SY (2009) Human hippocampal CA1 involvement during allocentric encoding of spatial information. *J Neurosci* 29:10512-10519.

- Suzuki K, Yamadori A, Hayakawa Y, Fujii T (1998) Pure Topographical Disorientation Related to Dysfunction of the Viewpoint Dependent Visual System. *Cortex* 34:589-599.
- Takahashi N, Kawamura M (2002) Pure Topographical Disorientation —The Anatomical Basis of Landmark Agnosia. *Cortex* 38:717-725.
- Takahashi N, Kawamura M, Shiota J, Kasahata N, Hirayama K (1997) Pure topographic disorientation due to right retrosplenial lesion. *Neurology* 49:464-469.
- Tamura I, Kitagawa M, Otsuki M, Kikuchi S, Tashiro K, Dubois B (2007) Pure topographical disorientation following a right forceps major of the splenium lesion: a case study. *Neurocase* 13:178-184.
- Taube JS (1995) Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *J Neurosci* 15:70-86.
- Taube JS, Muller RU, Ranck JB, Jr. (1990a) Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *J Neurosci* 10:436-447.
- Taube JS, Muller RU, Ranck JB, Jr. (1990b) Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J Neurosci* 10:420-435.
- Teng E, Squire LR (1999) Memory for places learned long ago is intact after hippocampal damage. *Nature* 400:675-677.
- Valiquette C, McNamara TP (2007) Different mental representations for place recognition and goal localization. *Psychonomic Bulletin & Review* 14:676-680.
- Van Leemput K, Bakkour A, Benner T, Wiggins G, Wald LL, Augustinack J, Dickerson BC, Golland P, Fischl B (2009) Automated segmentation of hippocampal subfields from ultra-high resolution in vivo MRI. *Hippocampus* 19:549-557.
- Vann SD, Aggleton JP, Maguire EA (2009) What does the retrosplenial cortex do? *Nature reviews Neuroscience* 10:792-802.
- Vass LK, Epstein RA (2013) Abstract representations of location and facing direction in the human brain. *J Neurosci* 33:6133-6142.
- Vinogradova OS (2001) Hippocampus as comparator: role of the two input and two output systems of the hippocampus in selection and registration of information. *Hippocampus* 11:578-598.
- Walther DB, Caddigan E, Fei-Fei L, Beck DM (2009) Natural scene categories revealed in distributed patterns of activity in the human brain. *J Neurosci* 29:10573-10581.
- Wang R, Spelke E (2002) Human spatial representation: insights from animals. *Trends Cogn Sci* 6:376.
- Wang RF (2012) Theories of spatial representations and reference frames: what can configuration errors tell us? *Psychon Bull Rev* 19:575-587.
- Whitlock JR, Sutherland RJ, Witter MP, Moser MB, Moser EI (2008) Navigating from hippocampus to parietal cortex. *Proc Natl Acad Sci U S A* 105:14755-14762.
- Whitlock JR, Pfuhl G, Dagslott N, Moser MB, Moser EI (2012) Functional split between parietal and entorhinal cortices in the rat. *Neuron* 73:789-802.
- Winter SS, Taube JS (2014) Head direction cells: from generation to integration. In: *Space, Time and Memory in the Hippocampal Formation* (Derdikman D, Knierim JJ, eds). New York: Springer.
- Wolbers T, Buchel C (2005) Dissociable retrosplenial and hippocampal contributions to successful formation of survey representations. *J Neurosci* 25:3333-3340.
- Wolbers T, Hegarty M (2010) What determines our navigational abilities? *Trends Cogn Sci* 14:138-146.
- Wolbers T, Wiener JM, Mallot HA, Buchel C (2007) Differential recruitment of the hippocampus, medial prefrontal cortex, and the human motion complex during path integration in humans. *J Neurosci* 27:9408-9416.
- Wolbers T, Hegarty M, Buchel C, Loomis JM (2008) Spatial updating: how the brain keeps track of changing object locations during observer motion. *Nat Neurosci* 11:1223-1230.

- Wolbers T, Klatzky RL, Loomis JM, Wutte MG, Giudice NA (2011) Modality-independent coding of spatial layout in the human brain. *Current biology : CB* 21:984-989.
- Yarkoni T, Poldrack RA, Nichols TE, Van Essen DC, Wager TD (2011) Large-scale automated synthesis of human functional neuroimaging data. *Nature methods* 8:665-670.
- Yassa MA, Stark CE (2009) A quantitative evaluation of cross-participant registration techniques for MRI studies of the medial temporal lobe. *Neuroimage* 44:319-327.
- Yoder RM, Clark BJ, Taube JS (2011) Origins of landmark encoding in the brain. *Trends in neurosciences* 34:561-571.
- Zhang H, Copara M, Ekstrom AD (2012) Differential recruitment of brain networks following route and cartographic map learning of spatial environments. *PloS one* 7:e44886.