THE ROLE OF THE RETROSPLENIAL CORTEX IN SPATIAL COGNITION

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

In Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Adam Michael Patrick Miller

December 2017

THE ROLE OF THE RETROSPLENIAL CORTEX IN SPATIAL COGNITION

Adam Michael Patrick Miller, Ph. D.

Cornell University 2017

Memory is a one of the most fundamental capacities of the brain. In addition to providing a repository for recent experiences, memory supports a wide range of cognitive functions including navigation, imagination, and sensory perception. While some facets of memory are well understood, including the essential role that the hippocampus plays in many forms of learning and the importance of neocortical regions for long-term memory storage, a complete mechanistic explanation for how memory supports cognition remains out of reach. A major obstacle is a lack of research into the form and development of neocortical memory representations, and how they participate in cognitive processes. To therefore provide insight into the role that neocortical regions play in memory, I examined the activity of large populations of neurons in the retrosplenial cortex (RSC) in awake, behaving rats as they performed spatial navigation tasks. The RSC is active during many memory-based cognitive functions in both humans and rats, and lesions of the RSC cause multiple forms of amnesia, including an inability to navigate. I find that neuronal activity patterns in the RSC form rich representations of two of the most important components of memory: space and time. I demonstrate that spatial representations develop slowly in the RSC over the course of many experiences, and that these representations are particularly sensitive to the spatial arrangement of the environment. I also found that the moment-to-moment nature of these representations correlates with the rat's

navigational proficiency, and that they can be used to predict the rat's future navigation behaviors, including upcoming memory failures. Lastly, I show that the activity of RSC ensembles selectively represents the rat's future goal location during instances when the rat must select between multiple destinations. Thus, my research provides support for the idea that neocortical memory regions support cognition through the reactivation of long-term memory representations relevant to the current behavior.

BIOGRAPHICAL SKETCH

Adam Miller was born in Stroudsburg, PA to Diane and Michael Miller. He has an older sister, Heather, and a younger brother, Scott. Adam attended Adlai E Stevenson High School in Lincolnshire, Illinois, before majoring in Psychology at Providence College in Providence, Rhode Island. At Providence, he assisted research on a variety of subjects, including obsessive-compulsive disorder, autonomic responses to stress in adolescence, and the relationship between personality types and learning styles. However, the majority of his work was conducted in the Providence Affective Neuroscience laboratory with Dr. Christopher Bloom testing an animal model of self-injury for its applicability to human forms of the condition. Adam joined Dr. David Smith in The Laboratory of Neurobiology of Learning and Memory at Cornell University in the fall of 2010, and began work on the role of the retrosplenial cortex in many forms of spatial cognition.

This dissertation is dedicated to my parents,

who instilled in me the values of independence, confidence, and creativity.

ACKNOWLEDGMENTS

I would first like to thank the many Smith Lab student researchers of my research group, Maze Team. Your work is on display throughout this dissertation. You helped collect the data day in and day out (or at least reminded me that I needed to go do that). You helped shape my ideas in discussions and journal clubs. By letting me help you with your writing and your presentations, you taught me how to write and how to present. Most importantly, you put up with all of my rambling stories and terrible jokes, and our conversations made lab time into time spent with friends. Thank you to Sarah Parauda, David Yu, Alex Tse, and Sherry Shi. And especially to my honors graduates, William Mau, Jason Li, and Anna Serrichio.

I would also like to thank the post-doctoral researchers that I had the pleasure of working with in our lab. You each embodied a unique approach to science that I strive to follow. Thank you Patrick Gil for your ebullient optimism and creativity. Thank you Dave Bulkin for your shrewd skepticism and criticism. Thank you Sasha Devore for your fierce brilliance and assiduousness. And a special thanks is due to Lindsey Vedder, who was a good friend and an important ally as we investigated similar questions.

Finally, I would like to thank my esteemed advisor, David Smith. David helped me to negotiate the many goals and interests that drive my research. He taught me to consider my own assumptions and beliefs, to pick and chose my starting points and trajectories for research projects. Most of all, he taught me what a good science project looks and feels like, and how to know when you have clearly conveyed a complex idea. Finally, and more practically, David provided an example for how to skillfully manage a group of scientists as they turn time and energy (and money) into data and knowledge. I hope to follow this example in the years ahead.

This research was generously supported by the National Institute of Mental Health grant number MH083809 awarded to David Smith.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF ABBREVIATIONS

RSC.....Retrosplenial cortex

LTM.....Long-term memory

ATN.....Anterior thalamic nuclei

RSA.....Retrosplenial dysgranular cortex

Rga.....Retrosplenial granular a cortex

Rgb.....Retrosplenial granular b cortex

fMRI.....Functional magnetic resonance imaging

BOLD.....Blood-oxygen-level dependent

IEG.....Immediate early gene

SEM.....Standard error of the mean

PREFACE

A major challenge in describing the study of cortical memory is expanding people's understanding of memory to include what they might more readily describe as "world knowledge." Memory is not like a recording that you can rewind in order to recount your recent sensations. It is also not like a bag, into which you might place facts or thoughts that you hope to re-experience later. Instead, I think of memory as a workable solution to an intractable problem: how can a brain that is trapped inside of a skull ever truly *know* the outside world?

Consider that, although we are surrounded by on all sides by people, places, and things, the brain doesn't truly "hear" or "see" any of them. Instead, the brain sits silently inside of its dark skull, receiving only electrical signals from different locations (ears, eyes, etc.) at different times. Insofar as these signals sound or look like anything at all, they have no resemblance to the real-world stuff that they represent: they are nothing more than electrical currents passing along wet cells. And, yet, we *do* have reliable knowledge of the world. We know what things are and where they are. And, to varying degrees, we know things about the past, and we know things about the future. But how?

The vast majority of this knowledge is acquired through our experience with the world. Some knowledge is gained rapidly, while other knowledge requires practice. The most basic forms of world knowledge improve your ability to sense the world. (For example, we can learn to think of horses and not zebras when we hear hoof beats, as the saying goes) More advanced knowledge about the underlying organization of the world allows us to predict, with some comfort, the vast majority of our future experiences.

How such knowledge forms, what it looks like in a physical brain, and this leads to the

many subjective, memory-related experiences that humans (and other animals) share, are urgent research questions in memory laboratories around the world. For me, these questions sit at the very heart of what it means to be human.

In the following chapters, I will introduce one small question: the role of one brain region, the retrosplenial cortex, in one type of memory processing, spatial representation. I will then present three chapters of original research on this question, first addressing the basic nature of this representation, and then describing how this representation forms and how it supports different kinds of behavior. Although this is only question, I believe that it is a central one, and I will argue in the final chapter that we can use the answers to this question to improve our understanding of memory more broadly.

CHAPTER ONE

Introduction

Navigating within a familiar environment requires a neural representation of the prominent features that define the environment (i.e., landmarks) and the spatial relationships among them. For example, traveling between two locations requires that you recognize your current position, select a destination from memory, and then remember how these two locations fit within the broader navigation context. Research conducted over the past 60 years has converged on the importance of interactions occurring over the course of learning between the limbic system and the neocortex for the formation and reactivation of these kinds of spatial representations in memory (i.e. a cognitive map; Calton & Taube, 2009; Moser, Kropff, & Moser, 2008; O'Keefe & Dostrovsky, 1971; Tolman, 1948; Vann, Aggleton, & Maguire, 2009). In particular, the hippocampus (of the limbic system) makes a prominent contribution early in learning by forming spatial representations of novel environments, and then by consolidating this information into neocortical regions such as the retrosplenial cortex (RSC) for long-term storage (Dudai, 2012; McClelland, McNaughton, & O'Reilly, 1995; R. G. Morris, 2006; Squire, 2004). After consolidation, neocortical representations are able to support navigation independent of the hippocampus (Day, Langston, & Morris, 2003; Maguire, Nannery, & Spiers, 2006; Tse et al., 2007; but see Winocur, Moscovitch, Caruana, & Binns, 2005; Winocur, Moscovitch, Fogel, Rosenbaum, & Sekeres, 2005), although the hippocampus continues to play an important role by reactivating relevant neocortical representations (Tanaka et al., 2014), and by consolidating new information (Tse et al., 2007; Tse et al., 2011). In this chapter, I will review evidence that the RSC plays an important role mediating interactions between limbic and neocortical regions, and

I will discuss its dual role in both long-term memory (LTM) storage and in the reactivation of long-term memory to support ongoing spatial cognition.

Functional Anatomy

The RSC is a posterior midline structure at the intersection between many limbic and cortical areas involved in spatial memory (**Figure 1)**. Historically, much attention has been paid to the connections between the RSC and the many limbic regions that are important for memory function in humans, such as the hippocampal formation and the anterior thalamic nuclei (ATN), while far less attention has been given to its connections to other neocortical regions such as the rhinal cortices or early visual areas. This led to an incomplete view of the RSC as simply another component of the limbic system (Papez, 1937), and also underestimated its roles in more traditionally-neocortical functions such as LTM storage and sensory perception. More modern descriptions of RSC connectivity highlight its dual role in both limbic and neocortical functions (e.g., Makino & Komiyama, 2015). Indeed, the RSC is often described as "intermediate" or "transition" cortex because of its position between the 3-layer archicortex of the hippocampus and the 6-layer neocortex (Vogt, 1976; Vogt & Laureys, 2005). This transition can be seen clearly in the cytoarchitecture of the RSC, with the granular region of the RSC (including the granular a region, Rga, also known as Broadmann's areas 29a and 29b, and the granular b region, Rgb, also known as Broadmann's area 29c) showing greater similarity to hippocampal archicortex and the dysgranular region of the RSC (RSA; area 30) showing greater similarity to neocortex (**Figure 2)**. Consistent with this cytoarchitecture, the granular RSC is preferentially connected with limbic regions while the dysgranular RSC is preferentially connected with neocortical regions (**Figure 3**; Burwell & Amaral, 1998; Insausti, Amaral, & Cowan, 1987;

Kobayashi & Amaral, 2003, 2007; Lavenex, Suzuki, & Amaral, 2002; R. Morris, Petrides, & Pandya, 1999; Suzuki & Amaral, 1994; van Groen & Wyss, 1990, 1992, 2003; Vogt, Pandya, & Rosene, 1987).

This pattern of connectivity places the RSC between limbic regions essential for memory formation and cortical regions essential for perception and LTM (Aggleton, 2008; Frankland & Bontempi, 2005; Maviel, Durkin, Menzaghi, & Bontempi, 2004). The importance of the RSC in supporting limbic function is well established, with the granular RSC forming a circuit with the ATN and the hippocampus that is essential for many kinds of memory. Damage to any of these three regions reliably disrupts functioning in the remaining two, including by eliminating synaptic plasticity or drastically reducing neural activity (Albasser, Poirier, Warburton, & Aggleton, 2007; Garden et al., 2009; Jenkins, Vann, Amin, & Aggleton, 2004; Poirier & Aggleton, 2009). These effects have predictably dire consequences for neural representations in these regions. For example, ATN lesions performed before training prevent the development of cue-elicited spiking activity in the RSC (Kubota & Gabriel, 1995; Smith, Freeman, Nicholson, & Gabriel, 2002), and lesions of the RSC impair the directional specificity of head direction cells in the ATN (Clark, Bassett, Wang, & Taube, 2010). ATN or RSC damage also disrupts spatial representations in the hippocampus: lesions of the ATN degrade the ability of hippocampal place fields to rotate with a shifted visual cue (Calton et al., 2003) and RSC inactivation during navigation causes hippocampal place fields to spontaneously shift to new locations (Cooper & Mizumori, 2001). Likewise, damage to the hippocampus impairs the normal development of cue-elicited neuronal responses in both the ATN and the RSC (Kang & Gabriel, 1998).

In contrast to our growing understanding of RSC-limbic circuit functionality, the functional importance of RSC-neocortical connections is far less understood. However, a major

clue comes from the broader functional organization of the neocortex, whereby the flow of sensory information from early sensory cortices, such as V1, to higher-level associative areas, such as the hippocampus, occurs along distinct processing streams. The dorsal stream, which connects V1 to the hippocampus though the parieto medial-temporal pathway, including the RSC, the posterior parietal cortex, and the parahippocampal region, plays an important role in the representation of spatial behavior ("where/how" information; (Goodale & Milner, 1992; Kravitz, Saleem, Baker, & Mishkin, 2011; Ungerleider & Mishkin, 1982), such as the coordination between visual inputs and movements (Perenin & Vighetto, 1988). In contrast, the ventral stream, which does not include the RSC, is essential for the representation of unique objects (Bussey & Saksida, 2007). To support the distributed representation of spatial and behavioral information, individual regions within the dorsal stream make distinct contributions ranging from the very specific (egocentric, local in time and place) in early sensory regions, to the very general (allocentric, distant in time and place), in later associative regions such as the hippocampus. Consistent with this, damage to regions along the dorsal stream has predictable effects depending on where it occurs (Aguirre & D'Esposito, 1999). For example, damage to the posterior parietal lobe (between the RSC and sensory cortex) produces more egocentric spatial-behavioral impairments, while damage to the parahippocampal cortex (between the RSC and the hippocampus) more commonly produces allocentric spatial-behavioral impairments. Damage to these regions also shows reliable differences in the nature of the resulting amnesia, with hippocampal-region damage producing especially profound anterograde amnesia, while damage closer to sensory cortex is more likely to result in retrograde amnesia as well. Retrograde amnesia observed following damage to these regions is likely due to their putative role in long-term memory storage (Danker & Anderson, 2010; Kellenbach, Brett, & Patterson, 2001), while

anterograde amnesia seen after damage to in other regions may result from the disruption of input to regions important for memory consolidation such as the hippocampus (e.g., Cooper & Mizumori, 2001).

Considerable attention has also been given to circuits (or loops) between cortical regions, the thalamus, and the basal ganglia, which are crucial for translating cortical representations into behavioral action. One perspective holds that cortical regions specify possible actions based on what the environment "affords," and these possible actions compete for selection such that only one is executed (Cisek, 2007; Cisek & Kalaska, 2010). In particular, cortical regions in the dorsal stream—in coordination with premotor regions—are believed to process sensory stimuli and to specify possible behaviors (Caligiore, Pezzulo, Miall, & Baldassarre, 2013). Although the RSC and the ATN are rarely included in descriptions of these loops (e.g., Yin & Knowlton, 2006), similar connections exist between these regions and the striatum (Aggleton et al., 2010; van Groen & Wyss, 1992, 2003), and interactions between these regions have been shown to influence behavioral selection and output (Gabriel, 1993). This suggests that spatial-contextual representations in the RSC may enable the specification of appropriate behaviors (e.g., navigation behaviors) that are later selected between via interactions with the ATN and striatum.

Spatial Representation

The above description of RSC connectivity suggests that the RSC might contribute to processes at the intersection between allocentric and egocentric spatial-behavioral representation. Studies of humans with RSC damage further describe these contributions by highlighting the role of the RSC in the representation of allocentric arrangements of navigational stimuli, such as landmarks. Damage to the RSC in humans occurs most frequently following cerebral

hemorrhages or tumors in the splenium of the corpus callosum. Patients with RSC damage consistently show navigation impairments, and cannot describe routes between locations or draw maps of the environment despite a preserved sense of familiarity and the ability to recognize individual landmarks and visual scenes (Maguire, 2001; **Figure 4**). In general, this navigation impairment ("topographical disorientation") can occur as a result of damage to any of the structures along the dorsal stream's parieto-medial temporal pathway (Aguirre & D'Esposito, 1999). In the case of RSC damage, topographical disorientation results from a compromised sense of allocentric space including an inability to learn the spatial relationships among landmarks (Obi, Bando, Takeda, & Sakuta, 1992; Takahashi, Kawamura, Shiota, Kasahata, & Hirayama, 1997). In particular, the RSC appears to support the ability to use navigation cues such as landmarks to discern the spatial layout of the environment (Maguire, 2001). For example, Ino and colleagues (Ino et al., 2007) describe the case of a patient who showed a striking impairment in the ability to use landmarks to navigate following a left RSC hemorrhage:

> *A 55 year-old right-handed man …who had been working as a taxicab driver in Kyoto City for 10 years, suddenly lost his knowledge of the route to his house while returning home after work. He could recognize buildings and the landscape and therefore understand where he was, but the landmarks that he recognized did not provoke directional information about any other places with respect to those landmarks. …He was unable to determine the direction to familiar destinations within the city with respect to his position at our hospital, although he had often visited this hospital while transporting passengers. He was also unable to describe or draw routes between his home and familiar places in his town, and could not draw the layout of his house. …He was unable to remember the route between his [hospital] room and the lavatory, which involved two right turns and one left turn, during several days of hospitalization.*

Although landmark cues can also support stimulus-response and egocentric navigation strategies, the disruption of cue processing in patients with RSC damage primarily impairs the

representation of allocentric space. RSC navigation impairments are closely related to the ability to localize within an environment, including the ability to place landmarks on a map (Takahashi et al., 1997) and the ability to locate one's current position in a miniature model of the testing room (Katayama, Takahashi, Ogawara, & Hattori, 1999). In contrast, RSC damage generally does not impair performance on traditional tests of egocentric representation, such as the ability to point to objects in the testing room from memory or the ability to recall the positions of landmarks that could be seen from the hospital window (Takahashi et al., 1997). Perhaps the strongest evidence for this distinction comes from one patient with RSC damage who was incapable of recalling the location of landmarks when questioned about their allocentric position, but who became capable of doing so when cued to adopt an egocentric strategy (i.e. his egocentric frame of reference was preserved; Takahashi et al., 1997, patient 2). When presented with a map of a familiar city but containing only streets and railroad tracks, the patient could not indicate the location of principal city buildings. However, when the examiner indicated a spot on the map and told the patient its name, the patient was able to recall the names and positions of buildings within sight of that location.

One possibility is that the RSC supports the use of landmarks in navigation by representing the allocentric arrangements of landmarks that define the spatial environment. Consistent with this, functional magnetic resonance imaging (fMRI) studies in humans have shown that the RSC is sensitive to changes in the layout of a visual scene (Dilks, Julian, Kubilius, Spelke, & Kanwisher, 2011), but is not sensitive to changes in the specific viewing angle (Morgan, Macevoy, Aguirre, & Epstein, 2011). Likewise, the RSC appears to participate more in spatial judgments related to landmarks than to simple recognition (Epstein, Parker, & Feiler, 2007). RSC BOLD activity is also closely related to the permanence of landmarks, a

feature that is essential to building a long-term representation of the layout of the environment that can support navigation (Auger & Maguire, 2013; Auger, Mullally, & Maguire, 2012; Auger, Zeidman, & Maguire, 2015a).

A growing body of work with rodents has also implicated the RSC in allocentric navigation. Lesions of the RSC impair performance on a wide range of navigation tasks, including the Morris water maze (Harker & Whishaw, 2004; Pothuizen, Davies, Aggleton, & Vann, 2010; Vann & Aggleton, 2002), radial arm maze (Keene & Bucci, 2009; Vann & Aggleton, 2002), and t-maze alternation tasks (Pothuizen et al., 2010). Rats with RSC damage are also impaired in representing landmark arrangements, and are incapable of detecting novel arrangements of objects, while remaining unimpaired at detecting novel objects themselves (Parron & Save, 2004; Vann & Aggleton, 2002). Likewise, early investigations of RSC neural firing have raised the possibility that RSC neurons exhibit the kinds of putatively allocentric firing patterns seen elsewhere in the hippocampal system, such as spatially localized firing (i.e. place cells; Chen, Lin, Barnes, & McNaughton, 1994; Chen, Lin, Green, Barnes, & McNaughton, 1994; Cho & Sharp, 2001). However, firing in these cells is often "sloppier" than is seen elsewhere. For example, RSC place fields are generally larger than their counterparts in the hippocampus (about three times as large in our studies; compare **Figure 5**, top and bottom; Smith, Barredo, & Mizumori, 2012), and RSC neurons have much higher background firing rates than hippocampal neurons, both inside and outside of their fields. This has led some researchers to question whether RSC neurons should be considered place cells at all (Alexander & Nitz, 2015).

Intriguingly, many RSC neurons respond to combinations of stimuli, including cues, locations, and/or directional heading—sometimes exhibiting surprising selectivity (Cho & Sharp,

2001). One RSC neuron described by Cho and Sharp (2001) fired preferentially when the rat was near the arena wall, turning left toward the center of the arena. This conjunctive coding is similar to that seen in hippocampal neurons, many of which respond selectively to combinations of locations, objects, and events – a mechanism thought to be a hallmark of relational memory functions (e.g., Komorowski, Manns, & Eichenbaum, 2009). Although these cells have not been thoroughly characterized in the RSC, it is possible that they encode the conjunctions of spatial position, navigation cues and the navigation behaviors associated with them. One of the clearest examples of this phenomenon comes from our laboratory, where we found that RSC neurons show unique responses upon receiving chocolate milk rewards depending on where on the maze they occurred (Smith et al., 2012).

Characterizing the representational properties of sloppy, combinatorial firing patterns, such as those seen in RSC neurons, requires statistical techniques that evaluate the entire population of recorded neurons simultaneously (as opposed to describing the firing properties of each neuron independently). This allows for agnostic descriptions of representational properties in terms of more practical relationships between distributed firing patterns and ongoing behavior, such as how well we can predict the current spatial location of the rat based on instantaneous neural firing. For example, one recent study evaluated the similarity of RSC population activity as rats performed left-and-right turns at multiple locations on a mobile maze that was moved between two locations in the testing room (Alexander & Nitz, 2015). The analysis focused on how well RSC firing encoded each of the three reference frames (turning behavior, rat-on-maze position, rat-in-room position), which vary in terms of their egocentric-allocentric qualities. Consistent with the above characterization of the RSC as between allocentric representation in limbic areas and more egocentric representation in cortical regions, this early application of

population-level statistical techniques to RSC neural activity found that each of the three reference frames were well encoded, perhaps suggesting that the RSC plays a translational role between the egocentric and allocentric. This work, however, did not address the more specific hypotheses of RSC representation outlined above, such as whether the RSC preferentially encodes landmark arrangements, or how predominant representations of each reference frame might support navigation.

Context Representation and Relational Coding

RSC contributions to spatial memory are likely just one example of a more general role in memory processing. For example, the RSC is also importantly involved in contextual memory, an ambiguous term referring to the many associations between co-occurring stimuli, such as environmental stimuli, behaviors, and expectations, that allow animals to tailor their behaviors for each environment they regularly encounter. Similar to the role of the RSC in spatial representation described above, lesions of the RSC impair the ability to associate a fearful outcome (electric shock) with the environmental stimuli that form a spatial context (Keene & Bucci, 2008a, 2008c). However, the contribution of the RSC appears to be more fundamental than simply recognizing or remembering spatial environments. Human studies using fMRI have shown that BOLD activity in the RSC increases in response to objects that are highly suggestive of particular contexts (e.g., a beach shell), or to arrays of objects that typically appear together within the same context (e.g., office items), even when the original spatial context is not presented (Bar & Aminoff, 2003). Most interestingly, the RSC also showed increased activation to arrays of objects that didn't share a spatial context at all, but were instead related by an abstract concept, such as "strength" or "birthday," suggesting that the RSC may simply be

encoding any reliable co-occurrences. Consistent with this, studies teasing apart RSC contributions to context processing have revealed that lesions of the RSC also impair the more basic ability to use compound stimuli, such as a tone-light combination, to guide behavior (Keene & Bucci, 2008b; Robinson, Keene, Iaccarino, Duan, & Bucci, 2011), as well as the ability to create passive associative links between stimuli simply because they co-occur (Robinson et al., 2011). Together, these findings suggest that the RSC appears to support context representations by processing prominent cues and the relationships between them. These findings may explain why some patients with RSC damage are impaired at paired-associate learning (Gainotti, Almonti, Di Betta, & Silveri, 1998), or associating items with the list in which they were presented (Bowers, Verfaellie, Valenstein, & Heilman, 1988).

Learning and Consolidation

Animals are not born with representations of their environment, but instead acquire them by exploring the environment over time. The acquisition of new information, and the process by which this information is then permanently stored (or at least maintained over long periods) is an area of active research. The classic view of this process, known as Standard Consolidation Theory (Marr, 1971; McClelland et al., 1995; Squire, 2004), holds that new sensory information moves through cortical regions to the hippocampus, where it is rapidly encoded through a process known as cellular consolidation. Over time, this information is then transferred back out of the hippocampus and into cortical regions for permanent LTM storage. Recent work on hippocampal-RSC interactions largely confirms this account. For example, rats performing an inhibitory avoidance task, where they learn to avoid an electrified area of the floor in an arena, require both their hippocampus and their RSC to remember the dangerous location in the first

days after learning (Katche, Dorman, Gonzalez, et al., 2013; Zanatta et al., 1997), but need only their RSC after about two weeks (Izquierdo & Medina, 1997; Katche, Dorman, Slipczuk, Cammarota, & Medina, 2013), suggesting that the memory was transferred into cortex. A closer look at this transfer from the hippocampus to the RSC revealed that an early wave of immediate early gene (IEG) activation occurs in the hippocampus within hours of the original learning experience—consistent with the cellular consolidation of this information into the hippocampus—that is then followed by simultaneous waves of IEG activation and protein synthesis in both the hippocampus and RSC around 12 hours after learning that are specifically necessary for the development of LTM (Katche et al., 2010; Katche, Dorman, Gonzalez, et al., 2013). Additional work on spatial context memory found that memory information consolidated into the RSC can then be reactivated by the hippocampus to support memory recall (Tanaka et al., 2014). Furthermore, stimulating the RSC directly, even if the hippocampus is inactivated, is sufficient to support memory recall, and also leads to the reactivation of other brain regions important for this kind of memory, such as the amygdala (Cowansage et al., 2014). This work confirms a growing body of research demonstrating that spatial and contextual memories continue to rely on the RSC after they have become hippocampal independent, leading a number of authors to suggest that critical elements of these memories are permanently stored in the RSC (Corcoran et al., 2011; Gusev & Gubin, 2010; Maviel et al., 2004).

Some evidence from fMRI studies of human subjects has also supported the idea that a hippocampal process drives the development of spatial representations in the RSC. Several fMRI studies have shown increases in BOLD activation in the RSC during spatial navigation (Aguirre, Detre, Alsop, & D'Esposito, 1996; Auger, Zeidman, & Maguire, 2015b; Gron, Wunderlich, Spitzer, Tomczak, & Riepe, 2000; Iaria, Chen, Guariglia, Ptito, & Petrides, 2007). In one such

study (Wolbers & Buchel, 2005; **Figure 6**), subjects received repeated first-person point-of-view tours of a virtual maze lined with visual landmarks. After each tour, participants were asked to imagine that they were facing a particular landmark and to indicate the relative position of a second landmark. With training, a subset of the participants became as accurate at judging the relative positions of landmarks located far apart on the maze as they were at judging pairs of landmarks encountered in sequence along the tour, suggesting that they had developed an allocentric representation of the maze. RSC activity increased as participants learned the task, consistent with its role in the development of an allocentric representation (**Figure 6B**). Interestingly, hippocampal activity did not increase over the course of training, but was instead correlated with the *slope* of the learning curve (**Figure 6C**) That is, the degree of hippocampal activation predicted the amount of learning occurring in a given session and, as a consequence, hippocampal activation declined with training as the subjects reached asymptote. This pattern of results is consistent with the idea that allocentric spatial representations develop slowly in the RSC, perhaps as a result of hippocampal-dependent processes supporting the incorporation of novel information into the developing RSC spatial representation.

Despite these findings that the RSC supports spatial LTM, some studies have also demonstrated a crucial role for the RSC in support of new learning, well before a hippocampal-cortical transfer could reasonably be expected. For example, the RSC is active during the learning stages of some virtual reality navigation tasks (Iaria et al., 2007), some functions in the RSC are necessary for learning stages of inhibitory avoidance learning (Katche, Dorman, Slipczuk, et al., 2013), and, at least in some cases, the RSC is necessary for the formation and the retrieval of recent context-fear memories (Corcoran et al., 2011; Keene & Bucci, 2008c). One plausible explanation of these findings is that RSC damage merely deprives the hippocampus of

the sensory input it needs to support learning (Cooper & Mizumori, 2001). However, a second—and more intriguing—possibility is that optimal performance in these tasks requires the use of memory information that was consolidated into the RSC before the task began. As a rule, all subjects enter a learning situation with relevant memories (no matter how basic) that either help or hinder performance, and some recent work has shown that rats with relevant spatial LTM can learn and consolidate new spatial information at an accelerated rate (Day et al., 2003; Tse et al., 2007; Tse et al., 2011). Consistent with the idea that the rapid learning resulted from the utilization of RSC LTM, this work showed that RSC IEG activation increased only when LTM facilitated new acquisition (Tse et al., 2011). Similarly, a study employing a classic chess paradigm for studying the contributions of LTM to new learning (Chase & Simon, 1973) found that the RSC was more sensitive to chess-typical arrangements than to random arrangements exclusively among expert chess players (Bilalic, Turella, Campitelli, Erb, & Grodd, 2012). Togther, these findings raise the possibility that learning impairments observed following RSC damage may be due to the brain's inability to access (and benefit from) general LTM information stored in the RSC.

Generative Memory

In addition to its role in learning, LTM also supports many other forms of cognitive and sensory processing, and the RSC is particularly well positioned to provide this support due to its connections with both limbic and sensory regions. Perhaps the best example comes from work on the interplay between the RSC and early visual areas over the course of learning (Makino & Komiyama, 2015; **Figure 7**). In this study, mice were trained to identify the direction of drift of a visual grating. Over the course of learning, the activity of V1 neurons became less influenced by

"bottom-up" input from the thalamus and became more influenced by "top-down" input from the RSC. When the researchers inactivated the RSC at the end of learning, they found that V1 neurons returned to their naïve (unlearned) state, and that the mice suddenly performed as though they had not received any training. Related phenomena have also been observed in humans, such as with the "boundary extension" phenomenon, a type of false memory whereby observers report having seen plausible information beyond the actual scene's boundaries (Intraub, Gottesman, Willey, & Zuk, 1996), which appears to draw on RSC processing (Park, Intraub, Yi, Widders, & Chun, 2007).

Additional evidence for an RSC role in generative memory processes comes from studies of the default network, a group of midline cortical structures, including the RSC, parahippocampal cortex, posterior parietal cortex, medial prefrontal cortex, and the hippocampus originally identified by their higher levels of BOLD activation during inter-trial rest periods than during goal-directed behavior (Raichle et al., 2001). The observation that these regions are often less active during task performance than during inter-trial rest has led some authors to describe them as comprising a "task-negative" network (Fox et al., 2005). Contrary to this, recent data suggest that default network activity reflects internally driven memory processing (Buckner, Andrews-Hanna, & Schacter, 2008; Mason et al., 2007; Spreng, 2012). Indeed, the default network is active during a range of episodic-like processes including self-projection (Buckner & Carroll, 2007), constructive memory (Hassabis & Maguire, 2007), and autobiographical memory (Spreng, Mar, & Kim, 2009). This network also supports the internal manipulation of memory items, such as when subjects combine items from long-term memory to construct imaginary mental scenes (Addis, Wong, & Schacter, 2007; Hassabis, Kumaran, & Maguire, 2007; Summerfield, Hassabis, & Maguire, 2010). Although the distributed nature of the default

network makes it unlikely that individuals with complete lesions of this network will be identified, naturally occurring conditions that compromise the default network, including Alzheimer's disease and normal ageing, are well known to impair performance on generative memory tasks (Buckner et al., 2005; Lustig et al., 2003).

Perhaps the best-known form of generative memory is episodic memory, which, like the other forms of memory processing described here, requires subjects to represent stimuli in space and time (Eichenbaum & Cohen, 2014; Zeidman & Maguire, 2016). Although apparently less common than spatial navigation deficits, RSC damage can also cause an amnesia for episodic memory that is similar to temporal lobe amnesia. The most famous case of RSC amnesia was described in patient T.R. who suffered a left RSC hemorrhage at the age of 39 (Bowers et al., 1988; Valenstein et al., 1987). T.R. experienced severe anterograde amnesia and a retrograde amnesia for events occurring in the 4 years preceding the injury, resulting in the loss of memory for the birth of his second child, a job change, and a recent relocation. Other patients with RSC lesions have shown even greater retrograde amnesia (e.g., ≥ 20 years in patient A.P.; Gainotti et al., 1998). Indeed, the similarity between RSC amnesia and other amnesic syndromes, combined with recent findings that damage to the hippocampus or ATN disrupts functioning in the RSC (Albasser et al., 2007; Garden et al., 2009; Jenkins et al., 2004; Poirier & Aggleton, 2009), has led some authors to suggest that memory impairments resulting from temporal lobe or diencephalic damage may be due, in part, to RSC dysfunction (Aggleton, 2008).

Working model of RSC-hippocampal interaction

The data reviewed here indicate that the RSC plays a critical role in spatial, contextual, and relational memory, and that interactions between the RSC and the hippocampus are essential

for normal memory processing. These data highlight the dual role of the RSC as both an input structure to the hippocampus that contributes importantly to the processing of cues related to navigation and context, and as a target of hippocampal output during the consolidation of this same information. It is likely that these two roles are closely related, with previously consolidated LTM augmenting the processing of feed-forward sensory input by the RSC (**Figure 8**). Specifically, we have presented evidence that (1) RSC input to the hippocampus is critical for the normal encoding of context representations. (2) Consolidation processes begin to establish memory traces in the cortex (McClelland et al., 1995), including the RSC, after encoding. In the fully trained subject, (3) the RSC provides input to the hippocampus that is crucial for identifying the active context, (4) thereby allowing the hippocampus to activate other relevant cortical representations and to incorporate novel input into the appropriate memory traces. Supporting these assertions, lesions of the RSC impair memory acquisition (Corcoran et al., 2011; Danker & Anderson, 2010; Ino et al., 2007; Katche, Dorman, Gonzalez, et al., 2013; Keene & Bucci, 2008c; Valenstein et al., 1987) as well as late-stage learning processes (Bussey, Muir, Everitt, & Robbins, 1996; Gabriel, 1993) and long-term memory (Danker & Anderson, 2010; Gainotti et al., 1998; Katche, Dorman, Gonzalez, et al., 2013; Takahashi et al., 1997; Valenstein et al., 1987). Additionally, the formation of long-term memory in the RSC depends on hippocampal-dependent processes (Katche, Dorman, Gonzalez, et al., 2013; Maviel et al., 2004; Staresina, Alink, Kriegeskorte, & Henson, 2014; Wolbers & Buchel, 2005), and RSC-hippocampal interactions continue to be important for learning in the fully trained subject, as RSC input supports hippocampal processing (Cooper & Mizumori, 2001), hippocampal activity can reactivate LTM representations in the RSC (Tanaka et al., 2014), and interplay between these regions facilitates the updating of existing memories (Tse et al., 2007; Tse et al., 2011).

The factors that influence the direction of information flow and the interplay between new learning, consolidation, and memory updating in the RSC and hippocampus remain poorly understood and future studies will need to monitor both regions simultaneously in order to formulate a truly detailed account.

Concluding remarks

Investigations of spatial navigation provide a critical window into the mechanisms of memory representation throughout the extended hippocampal system. The RSC is an important focus of memory research due to its intimate interconnectivity with many limbic and cortical regions involved in spatial representation, as well as its dual role in feedforward sensory processing and long-term memory storage. The RSC contribution to spatial representation is closely related to a more general role representing the long-term associations that comprise context representations. Such long-term associations in the RSC appear to develop in collaboration with the hippocampus, which supports the consolidation of new spatial information and the updating of existing cortical representations. Although these data present an optimistic picture of our understanding of RSC processing, many questions remain, especially concerning the nature of RSC long-term representations, how they form, how they are updated, and how they are utilized in support of ongoing processing.

References

Addis, D. R., Wong, A. T., & Schacter, D. L. (2007). Remembering the past and imagining the future: common and distinct neural substrates during event construction and elaboration. *Neuropsychologia, 45*(7), 1363-1377. doi: 10.1016/j.neuropsychologia.2006.10.016

Aggleton, J. P. (2008). EPS Mid-Career Award 2006. Understanding anterograde amnesia: disconnections and hidden lesions. *Q J Exp Psychol (Hove), 61*(10), 1441-1471. doi: 10.1080/17470210802215335

Aggleton, J. P., O'Mara, S. M., Vann, S. D., Wright, N. F., Tsanov, M., & Erichsen, J. T. (2010). Hippocampal-anterior thalamic pathways for memory: uncovering a network of direct and indirect actions. *Eur J Neurosci, 31*(12), 2292-2307. doi: 10.1111/j.1460-9568.2010.07251.x

Aguirre, G. K., & D'Esposito, M. (1999). Topographical disorientation: a synthesis and taxonomy. *Brain, 122 (Pt 9)*, 1613-1628.

Aguirre, G. K., Detre, J. A., Alsop, D. C., & D'Esposito, M. (1996). The parahippocampus subserves topographical learning in man. *Cereb Cortex, 6*(6), 823-829.

Albasser, M. M., Poirier, G. L., Warburton, E. C., & Aggleton, J. P. (2007). Hippocampal lesions halve immediate-early gene protein counts in retrosplenial cortex: distal dysfunctions in a spatial memory system. *Eur J Neurosci, 26*(5), 1254-1266. doi: 10.1111/j.1460-9568.2007.05753.x

Alexander, A. S., & Nitz, D. A. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nat Neurosci, 18*(8), 1143-1151. doi: 10.1038/nn.4058

Auger, S. D., & Maguire, E. A. (2013). Assessing the mechanism of response in the retrosplenial cortex of good and poor navigators. *Cortex, 49*(10), 2904-2913. doi: 10.1016/j.cortex.2013.08.002

Auger, S. D., Mullally, S. L., & Maguire, E. A. (2012). Retrosplenial cortex codes for permanent landmarks. *PLoS One, 7*(8), e43620. doi: 10.1371/journal.pone.0043620

Auger, S. D., Zeidman, P., & Maguire, E. A. (2015a). A central role for the retrosplenial cortex in de novo environmental learning. *eLife, 4*. doi: 10.7554/eLife.09031

Auger, S. D., Zeidman, P., & Maguire, E. A. (2015b). A central role for the retrosplenial cortex in de novo environmental learning. *eLife, 4*, e09031. doi: 10.7554/eLife.09031

Bar, M., & Aminoff, E. (2003). Cortical analysis of visual context. *Neuron, 38*(2), 347-358.

Bilalic, M., Turella, L., Campitelli, G., Erb, M., & Grodd, W. (2012). Expertise modulates the neural basis of context dependent recognition of objects and their relations. *Hum Brain Mapp, 33*(11), 2728-2740. doi: 10.1002/hbm.21396

Bowers, D., Verfaellie, M., Valenstein, E., & Heilman, K. M. (1988). Impaired acquisition of temporal information in retrosplenial amnesia. *Brain Cogn, 8*(1), 47-66.

Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network: anatomy, function, and relevance to disease. *Ann N Y Acad Sci, 1124*, 1-38. doi: 10.1196/annals.1440.011

Buckner, R. L., & Carroll, D. C. (2007). Self-projection and the brain. *Trends Cogn Sci, 11*(2), 49-57. doi: 10.1016/j.tics.2006.11.004

Buckner, R. L., Snyder, A. Z., Shannon, B. J., LaRossa, G., Sachs, R., Fotenos, A. F., . . . Mintun, M. A. (2005). Molecular, structural, and functional characterization of Alzheimer's disease: evidence for a relationship between default activity, amyloid, and memory. *J Neurosci, 25*(34), 7709-7717. doi: 10.1523/jneurosci.2177-05.2005

Burwell, R. D., & Amaral, D. G. (1998). Cortical afferents of the perirhinal, postrhinal, and entorhinal cortices of the rat. *J Comp Neurol, 398*(2), 179-205.

Bussey, T. J., Muir, J. L., Everitt, B. J., & Robbins, T. W. (1996). Dissociable effects of anterior and posterior cingulate cortex lesions on the acquisition of a conditional visual discrimination: facilitation of early learning vs. impairment of late learning. *Behav Brain Res, 82*(1), 45-56.

Bussey, T. J., & Saksida, L. M. (2007). Memory, perception, and the ventral visual-perirhinal-hippocampal stream: thinking outside of the boxes. *Hippocampus, 17*(9), 898-908. doi: 10.1002/hipo.20320

Caligiore, D., Pezzulo, G., Miall, R. C., & Baldassarre, G. (2013). The contribution of brain sub-cortical loops in the expression and acquisition of action understanding abilities. *Neurosci Biobehav Rev, 37*(10 Pt 2), 2504-2515. doi: 10.1016/j.neubiorev.2013.07.016

Calton, J. L., Stackman, R. W., Goodridge, J. P., Archey, W. B., Dudchenko, P. A., & Taube, J. S. (2003). Hippocampal place cell instability after lesions of the head direction cell network. *J Neurosci, 23*(30), 9719-9731.

Calton, J. L., & Taube, J. S. (2009). Where am I and how will I get there from here? A role for posterior parietal cortex in the integration of spatial information and route planning. *Neurobiol Learn Mem, 91*(2), 186-196. doi: 10.1016/j.nlm.2008.09.015

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cogn Psychol, 4*(1), 55-81. doi: 10.1016/0010-0285(73)90004-2

Chen, L. L., Lin, L. H., Barnes, C. A., & McNaughton, B. L. (1994). Head-direction cells in the rat posterior cortex. II. Contributions of visual and ideothetic information to the directional firing. *Exp Brain Res, 101*(1), 24-34.

Chen, L. L., Lin, L. H., Green, E. J., Barnes, C. A., & McNaughton, B. L. (1994). Head-direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. *Exp Brain Res, 101*(1), 8-23.

Cho, J., & Sharp, P. E. (2001). Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behav Neurosci, 115*(1), 3-25.

Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philos Trans R Soc Lond B Biol Sci, 362*(1485), 1585-1599. doi: 10.1098/rstb.2007.2054

Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annu Rev Neurosci, 33*, 269-298. doi: 10.1146/annurev.neuro.051508.135409

Clark, B. J., Bassett, J. P., Wang, S. S., & Taube, J. S. (2010). Impaired head direction cell representation in the anterodorsal thalamus after lesions of the retrosplenial cortex. *J Neurosci, 30*(15), 5289-5302. doi: 10.1523/jneurosci.3380-09.2010

Cooper, B. G., & Mizumori, S. J. (2001). Temporary inactivation of the retrosplenial cortex causes a transient reorganization of spatial coding in the hippocampus. *J Neurosci, 21*(11), 3986-4001.

Corcoran, K. A., Donnan, M. D., Tronson, N. C., Guzman, Y. F., Gao, C., Jovasevic, V., . . . Radulovic, J. (2011). NMDA receptors in retrosplenial cortex are necessary for retrieval of recent and remote context fear memory. *J Neurosci, 31*(32), 11655-11659. doi: 10.1523/jneurosci.2107-11.2011

Cowansage, Kiriana K., Shuman, T., Dillingham, Blythe C., Chang, A., Golshani, P., & Mayford, M. (2014). Direct Reactivation of a Coherent Neocortical Memory of Context. *Neuron, 84*(2), 432-441. doi: 10.1016/j.neuron.2014.09.022

Danker, J. F., & Anderson, J. R. (2010). The ghosts of brain states past: remembering reactivates the brain regions engaged during encoding. *Psychol Bull, 136*(1), 87-102. doi: 10.1037/a0017937

Day, M., Langston, R., & Morris, R. G. (2003). Glutamate-receptor-mediated encoding and retrieval of paired-associate learning. *Nature, 424*(6945), 205-209. doi: 10.1038/nature01769

Dilks, D. D., Julian, J. B., Kubilius, J., Spelke, E. S., & Kanwisher, N. (2011). Mirror-image sensitivity and invariance in object and scene processing pathways. *J Neurosci, 31*(31), 11305-11312. doi: 10.1523/jneurosci.1935-11.2011

Dudai, Y. (2012). The restless engram: consolidations never end. *Annu Rev Neurosci, 35*, 227-247. doi: 10.1146/annurev-neuro-062111-150500

Eichenbaum, H., & Cohen, Neal J. (2014). Can We Reconcile the Declarative Memory and Spatial Navigation Views on Hippocampal Function? *Neuron, 83*(4), 764-770. doi: 10.1016/j.neuron.2014.07.032

Epstein, R. A., Parker, W. E., & Feiler, A. M. (2007). Where am I now? Distinct roles for parahippocampal and retrosplenial cortices in place recognition. *J Neurosci, 27*(23), 6141-6149. doi: 10.1523/jneurosci.0799-07.2007

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proc Natl Acad Sci U S A, 102*(27), 9673-9678. doi: 10.1073/pnas.0504136102

Frankland, P. W., & Bontempi, B. (2005). The organization of recent and remote memories. *Nat Rev Neurosci, 6*(2), 119-130. doi: 10.1038/nrn1607

Gabriel, M. (1993). Discriminative avoidance learning: a model system *Neurobiology of cingulate cortex and limbic thalamus* (pp. 478-523). Boston: Birkhäuser.

Gainotti, G., Almonti, S., Di Betta, A. M., & Silveri, M. C. (1998). Retrograde amnesia in a patient with retrosplenial tumour. *Neurocase, 4*(6), 519-526. doi: 10.1080/13554799808410644

Garden, D. L., Massey, P. V., Caruana, D. A., Johnson, B., Warburton, E. C., Aggleton, J. P., & Bashir, Z. I. (2009). Anterior thalamic lesions stop synaptic plasticity in retrosplenial cortex slices: expanding the pathology of diencephalic amnesia. *Brain, 132*(Pt 7), 1847-1857. doi: 10.1093/brain/awp090

Goodale, M. A., & Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends Neurosci, 15*(1), 20-25.

Gron, G., Wunderlich, A. P., Spitzer, M., Tomczak, R., & Riepe, M. W. (2000). Brain activation during human navigation: gender-different neural networks as substrate of performance. *Nat Neurosci, 3*(4), 404-408. doi: 10.1038/73980

Gusev, P. A., & Gubin, A. N. (2010). Arc/Arg3.1 mRNA global expression patterns elicited by memory recall in cerebral cortex differ for remote versus recent spatial memories. *Front Integr Neurosci, 4*, 15. doi: 10.3389/fnint.2010.00015

Harker, K. T., & Whishaw, I. Q. (2004). Impaired place navigation in place and matching-to-place swimming pool tasks follows both retrosplenial cortex lesions and cingulum bundle lesions in rats. *Hippocampus, 14*(2), 224-231. doi: 10.1002/hipo.10159

Hassabis, D., Kumaran, D., & Maguire, E. A. (2007). Using imagination to understand the neural basis of episodic memory. *J Neurosci, 27*(52), 14365-14374. doi: 10.1523/jneurosci.4549-07.2007

Hassabis, D., & Maguire, E. A. (2007). Deconstructing episodic memory with construction. *Trends Cogn Sci, 11*(7), 299-306. doi: 10.1016/j.tics.2007.05.001

Iaria, G., Chen, J. K., Guariglia, C., Ptito, A., & Petrides, M. (2007). Retrosplenial and hippocampal brain regions in human navigation: complementary functional contributions to the formation and use of cognitive maps. *Eur J Neurosci, 25*(3), 890-899. doi: 10.1111/j.1460-9568.2007.05371.x

Ino, T., Doi, T., Hirose, S., Kimura, T., Ito, J., & Fukuyama, H. (2007). Directional disorientation following left retrosplenial hemorrhage: a case report with fMRI studies. *Cortex, 43*(2), 248-254.

Insausti, R., Amaral, D. G., & Cowan, W. M. (1987). The entorhinal cortex of the monkey: II. Cortical afferents. *J Comp Neurol, 264*(3), 356-395. doi: 10.1002/cne.902640306

Intraub, H., Gottesman, C. V., Willey, E. V., & Zuk, I. J. (1996). Boundary Extension for Briefly Glimpsed Photographs: Do Common Perceptual Processes Result in Unexpected Memory Distortions? *Journal of Memory and Language, 35*(2), 118-134. doi: 10.1006/jmla.1996.0007

Izquierdo, I., & Medina, J. H. (1997). Memory formation: the sequence of biochemical events in the hippocampus and its connection to activity in other brain structures. *Neurobiol Learn Mem, 68*(3), 285-316. doi: 10.1006/nlme.1997.3799

Jenkins, T. A., Vann, S. D., Amin, E., & Aggleton, J. P. (2004). Anterior thalamic lesions stop immediate early gene activation in selective laminae of the retrosplenial cortex: evidence of covert pathology in rats? *Eur J Neurosci, 19*(12), 3291-3304. doi: 10.1111/j.0953-816X.2004.03421.x

Kang, E., & Gabriel, M. (1998). Hippocampal modulation of cingulo-thalamic neuronal activity and discriminative avoidance learning in rabbits. *Hippocampus, 8*(5), 491-510. doi: 10.1002/(SICI)1098-1063(1998)8:5<491::AID-HIPO8>3.0.CO;2-C

Katayama, K., Takahashi, N., Ogawara, K., & Hattori, T. (1999). Pure topographical disorientation due to right posterior cingulate lesion. *Cortex, 35*(2), 279-282.

Katche, C., Bekinschtein, P., Slipczuk, L., Goldin, A., Izquierdo, I. A., Cammarota, M., & Medina, J. H. (2010). Delayed wave of c-Fos expression in the dorsal hippocampus involved specifically in persistence of long-term memory storage. *Proc Natl Acad Sci U S A, 107*(1), 349-354. doi: 10.1073/pnas.0912931107

Katche, C., Dorman, G., Gonzalez, C., Kramar, C. P., Slipczuk, L., Rossato, J. I., . . . Medina, J. H. (2013). On the role of retrosplenial cortex in long-lasting memory storage. *Hippocampus, 23*(4), 295-302. doi: 10.1002/hipo.22092

Katche, C., Dorman, G., Slipczuk, L., Cammarota, M., & Medina, J. H. (2013). Functional integrity of the retrosplenial cortex is essential for rapid consolidation and recall of fear memory. *Learn Mem, 20*(4), 170-173. doi: 10.1101/lm.030080.112

Keene, C. S., & Bucci, D. J. (2008a). Contributions of the retrosplenial and posterior parietal cortices to cue-specific and contextual fear conditioning. *Behav Neurosci, 122*(1), 89-97. doi: 10.1037/0735-7044.122.1.89

Keene, C. S., & Bucci, D. J. (2008b). Involvement of the retrosplenial cortex in processing multiple conditioned stimuli. *Behav Neurosci, 122*(3), 651-658. doi: 10.1037/0735-7044.122.3.651

Keene, C. S., & Bucci, D. J. (2008c). Neurotoxic lesions of retrosplenial cortex disrupt signaled and unsignaled contextual fear conditioning. *Behav Neurosci, 122*(5), 1070-1077. doi: 10.1037/a0012895

Keene, C. S., & Bucci, D. J. (2009). Damage to the retrosplenial cortex produces specific impairments in spatial working memory. *Neurobiol Learn Mem, 91*(4), 408-414. doi: 10.1016/j.nlm.2008.10.009

Kellenbach, M. L., Brett, M., & Patterson, K. (2001). Large, colorful, or noisy? Attribute- and modality-specific activations during retrieval of perceptual attribute knowledge. *Cogn Affect Behav Neurosci, 1*(3), 207-221.

Kobayashi, Y., & Amaral, D. G. (2000). Macaque monkey retrosplenial cortex: I. three-dimensional and cytoarchitectonic organization. *J Comp Neurol, 426*(3), 339-365.

Kobayashi, Y., & Amaral, D. G. (2003). Macaque monkey retrosplenial cortex: II. Cortical afferents. *J Comp Neurol, 466*(1), 48-79. doi: 10.1002/cne.10883

Kobayashi, Y., & Amaral, D. G. (2007). Macaque monkey retrosplenial cortex: III. Cortical efferents. *J Comp Neurol, 502*(5), 810-833. doi: 10.1002/cne.21346

Komorowski, R. W., Manns, J. R., & Eichenbaum, H. (2009). Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. *J Neurosci, 29*(31), 9918-9929. doi: 10.1523/jneurosci.1378-09.2009

Kravitz, D. J., Saleem, K. S., Baker, C. I., & Mishkin, M. (2011). A new neural framework for visuospatial processing. *Nat Rev Neurosci, 12*(4), 217-230. doi: 10.1038/nrn3008

Kubota, Y., & Gabriel, M. (1995). Studies of the limbic comparator: limbic circuit training-induced unit activity and avoidance behavior in rabbits with anterior dorsal thalamic lesions. *Behav Neurosci, 109*(2), 258-277.

Lavenex, P., Suzuki, W. A., & Amaral, D. G. (2002). Perirhinal and parahippocampal cortices of the macaque monkey: projections to the neocortex. *J Comp Neurol, 447*(4), 394-420. doi: 10.1002/cne.10243

Lustig, C., Snyder, A. Z., Bhakta, M., O'Brien, K. C., McAvoy, M., Raichle, M. E., . . . Buckner, R. L. (2003). Functional deactivations: change with age and dementia of the Alzheimer type. *Proc Natl Acad Sci U S A, 100*(24), 14504-14509. doi: 10.1073/pnas.2235925100

Maguire, E. A. (2001). The retrosplenial contribution to human navigation: a review of lesion and neuroimaging findings. *Scand J Psychol, 42*(3), 225-238.

Maguire, E. A., Nannery, R., & Spiers, H. J. (2006). Navigation around London by a taxi driver with bilateral hippocampal lesions. *Brain, 129*(Pt 11), 2894-2907. doi: 10.1093/brain/awl286

Makino, H., & Komiyama, T. (2015). Learning enhances the relative impact of top-down processing in the visual cortex. *Nat Neurosci, 18*(8), 1116-1122. doi: 10.1038/nn.4061

Marr, D. (1971). Simple Memory: A Theory for Archicortex. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 262*(841), 23-81.

Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T., & Macrae, C. N. (2007). Wandering minds: the default network and stimulus-independent thought. *Science, 315*(5810), 393-395. doi: 10.1126/science.1131295

Maviel, T., Durkin, T. P., Menzaghi, F., & Bontempi, B. (2004). Sites of neocortical reorganization critical for remote spatial memory. *Science, 305*(5680), 96-99. doi: 10.1126/science.1098180

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev, 102*(3), 419-457.

Morgan, L. K., Macevoy, S. P., Aguirre, G. K., & Epstein, R. A. (2011). Distances between real-world locations are represented in the human hippocampus. *J Neurosci, 31*(4), 1238-1245. doi: 10.1523/jneurosci.4667-10.2011

Morris, R., Petrides, M., & Pandya, D. N. (1999). Architecture and connections of retrosplenial area 30 in the rhesus monkey (Macaca mulatta). *Eur J Neurosci, 11*(7), 2506-2518.

Morris, R. G. (2006). Elements of a neurobiological theory of hippocampal function: the role of synaptic plasticity, synaptic tagging and schemas. *Eur J Neurosci, 23*(11), 2829-2846. doi: 10.1111/j.1460-9568.2006.04888.x

Moser, E. I., Kropff, E., & Moser, M. B. (2008). Place cells, grid cells, and the brain's spatial representation system. *Annu Rev Neurosci, 31*, 69-89. doi: 10.1146/annurev.neuro.31.061307.090723

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res, 34*(1), 171-175.

Obi, T., Bando, M., Takeda, K., & Sakuta, M. (1992). [A case of topographical disturbance following a left medial parieto-occipital lobe infarction]. *Rinsho Shinkeigaku, 32*(4), 426-429.

Papez. (1937). A proposed mechanism of emotion. 1937 [classical article]. *The Journal of Neuropsychiatry and Clinical Neurosciences, 7*(1), 103-112. doi: 10.1176/jnp.7.1.103

Park, S., Intraub, H., Yi, D. J., Widders, D., & Chun, M. M. (2007). Beyond the edges of a view: boundary extension in human scene-selective visual cortex. *Neuron, 54*(2), 335-342. doi: 10.1016/j.neuron.2007.04.006

Parron, C., & Save, E. (2004). Comparison of the effects of entorhinal and retrosplenial cortical lesions on habituation, reaction to spatial and non-spatial changes during object exploration in the rat. *Neurobiol Learn Mem, 82*(1), 1-11. doi: 10.1016/j.nlm.2004.03.004

Perenin, M. T., & Vighetto, A. (1988). Optic ataxia: a specific disruption in visuomotor mechanisms. I. Different aspects of the deficit in reaching for objects. *Brain, 111 (Pt 3)*, 643-674.

Poirier, G. L., & Aggleton, J. P. (2009). Post-surgical interval and lesion location within the limbic thalamus determine extent of retrosplenial cortex immediate-early gene hypoactivity. *Neuroscience, 160*(2), 452-469. doi: 10.1016/j.neuroscience.2009.02.021

Pothuizen, H. H., Davies, M., Aggleton, J. P., & Vann, S. D. (2010). Effects of selective granular retrosplenial cortex lesions on spatial working memory in rats. *Behav Brain Res, 208*(2), 566-575. doi: 10.1016/j.bbr.2010.01.001

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proc Natl Acad Sci U S A, 98*(2), 676-682. doi: 10.1073/pnas.98.2.676

Robinson, S., Keene, C. S., Iaccarino, H. F., Duan, D., & Bucci, D. J. (2011). Involvement of retrosplenial cortex in forming associations between multiple sensory stimuli. *Behav Neurosci, 125*(4), 578-587. doi: 10.1037/a0024262

Smith, D. M., Barredo, J., & Mizumori, S. J. (2012). Complimentary roles of the hippocampus and retrosplenial cortex in behavioral context discrimination. *Hippocampus, 22*(5), 1121-1133. doi: 10.1002/hipo.20958

Smith, D. M., Freeman, J. H., Jr., Nicholson, D., & Gabriel, M. (2002). Limbic thalamic lesions, appetitively motivated discrimination learning, and training-induced neuronal activity in rabbits. *J Neurosci, 22*(18), 8212-8221.

Spreng, R. N. (2012). The fallacy of a "task-negative" network. *Front Psychol, 3*, 145. doi: 10.3389/fpsyg.2012.00145

Spreng, R. N., Mar, R. A., & Kim, A. S. (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*(3), 489-510. doi: 10.1162/jocn.2008.21029

Squire, L. R. (2004). Memory systems of the brain: a brief history and current perspective. *Neurobiol Learn Mem, 82*(3), 171-177. doi: 10.1016/j.nlm.2004.06.005

Staresina, B. P., Alink, A., Kriegeskorte, N., & Henson, R. N. (2014). Awake reactivation predicts memory in humans. *Proc Natl Acad Sci U S A, 110*(52), 21159-21164. doi: 10.1073/pnas.1311989110

Summerfield, J. J., Hassabis, D., & Maguire, E. A. (2010). Differential engagement of brain regions within a 'core' network during scene construction. *Neuropsychologia, 48*(5), 1501-1509. doi: 10.1016/j.neuropsychologia.2010.01.022

Suzuki, W. A., & Amaral, D. G. (1994). Topographic organization of the reciprocal connections between the monkey entorhinal cortex and the perirhinal and parahippocampal cortices. *J Neurosci, 14*(3 Pt 2), 1856-1877.

Takahashi, N., Kawamura, M., Shiota, J., Kasahata, N., & Hirayama, K. (1997). Pure topographic disorientation due to right retrosplenial lesion. *Neurology, 49*(2), 464-469.

Tanaka, Kazumasa Z., Pevzner, A., Hamidi, Anahita B., Nakazawa, Y., Graham, J., & Wiltgen, Brian J. (2014). Cortical Representations Are Reinstated by the Hippocampus during Memory Retrieval. *Neuron, 84*(2), 347-354. doi: 10.1016/j.neuron.2014.09.037

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychol Rev, 55*(4), 189-208.

Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., . . . Morris, R. G. (2007). Schemas and memory consolidation. *Science, 316*(5821), 76-82. doi: 10.1126/science.1135935

Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., . . . Morris, R. G. (2011). Schema-dependent gene activation and memory encoding in neocortex. *Science, 333*(6044), 891-895. doi: 10.1126/science.1205274

Ungerleider, L. G., & Mishkin, M. (1982). Two Cortical Visual Systems. In D. J. Ingle, M. Goodale & R. J. W. Mansfield (Eds.), *Analysis of Visual Behaviour* (pp. 549-586).

Valenstein, E., Bowers, D., Verfaellie, M., Heilman, K. M., Day, A., & Watson, R. T. (1987). Retrosplenial amnesia. *Brain, 110 (Pt 6)*, 1631-1646.

van Groen, T., & Wyss, J. M. (1990). Connections of the retrosplenial granular a cortex in the rat. *J Comp Neurol, 300*(4), 593-606. doi: 10.1002/cne.903000412

van Groen, T., & Wyss, J. M. (1992). Connections of the retrosplenial dysgranular cortex in the rat. *J Comp Neurol, 315*(2), 200-216. doi: 10.1002/cne.903150207

van Groen, T., & Wyss, J. M. (2003). Connections of the retrosplenial granular b cortex in the rat. *J Comp Neurol, 463*(3), 249-263. doi: 10.1002/cne.10757

Vann, S. D., & Aggleton, J. P. (2002). Extensive cytotoxic lesions of the rat retrosplenial cortex reveal consistent deficits on tasks that tax allocentric spatial memory. *Behav Neurosci, 116*(1), 85-94.

Vann, S. D., Aggleton, J. P., & Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nat Rev Neurosci, 10*(11), 792-802. doi: 10.1038/nrn2733

Vogt, B. A. (1976). Retrosplenial cortex in the rhesus monkey: a cytoarchitectonic and Golgi study. *J Comp Neurol, 169*(1), 63-97. doi: 10.1002/cne.901690105

Vogt, B. A., & Laureys, S. (2005). Posterior cingulate, precuneal and retrosplenial cortices: cytology and components of the neural network correlates of consciousness. *Prog Brain Res, 150*, 205-217. doi: 10.1016/s0079-6123(05)50015-3

Vogt, B. A., Pandya, D. N., & Rosene, D. L. (1987). Cingulate cortex of the rhesus monkey: I. Cytoarchitecture and thalamic afferents. *J Comp Neurol, 262*(2), 256-270. doi: 10.1002/cne.902620207

Winocur, G., Moscovitch, M., Caruana, D. A., & Binns, M. A. (2005). Retrograde amnesia in rats with lesions to the hippocampus on a test of spatial memory. *Neuropsychologia, 43*(11), 1580-1590. doi: 10.1016/j.neuropsychologia.2005.01.013

Winocur, G., Moscovitch, M., Fogel, S., Rosenbaum, R. S., & Sekeres, M. (2005). Preserved spatial memory after hippocampal lesions: effects of extensive experience in a complex environment. *Nat Neurosci, 8*(3), 273-275. doi: 10.1038/nn1401

Wolbers, T., & Buchel, C. (2005). Dissociable retrosplenial and hippocampal contributions to successful formation of survey representations. *J Neurosci, 25*(13), 3333-3340. doi: 10.1523/jneurosci.4705-04.2005

Yin, H. H., & Knowlton, B. J. (2006). The role of the basal ganglia in habit formation. *Nat Rev Neurosci, 7*(6), 464-476. doi: 10.1038/nrn1919

Zanatta, M. S., Quillfeldt, J. H., Schaeffer, E., Schmitz, P. K., Quevedo, J., Medina, J. H., & Izquierdo, I. (1997). Involvement of the hippocampus, amygdala, entorhinal cortex and posterior parietal cortex in memory consolidation. *Braz J Med Biol Res, 30*(2), 235-240.

Zeidman, P., & Maguire, E. A. (2016). Anterior hippocampus: the anatomy of perception, imagination and episodic memory. *Nat Rev Neurosci, 17*(3), 173-182. doi: 10.1038/nrn.2015.24

Figures

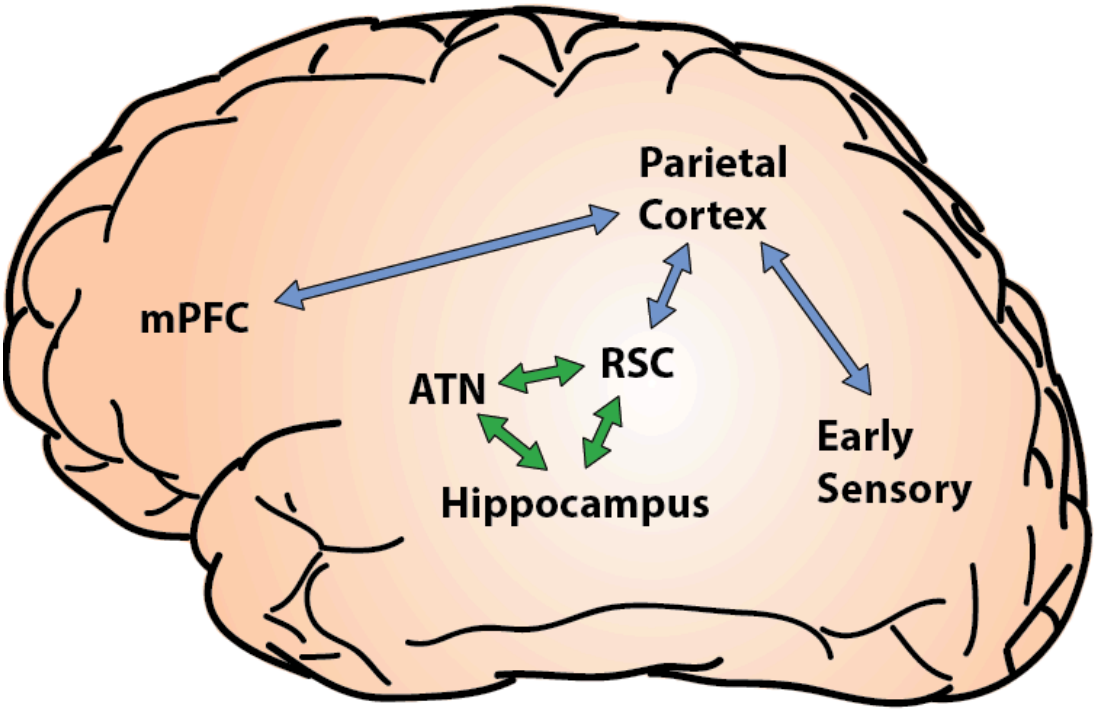

Figure 1 Overview of RSC connectivity. The RSC is centrally positioned between cortical regions (blue) and limbic regions (green). The parietal lobe merges information arriving from early sensory areas, and shares connectivity with both the medial prefrontal cortex (mPFC) and the RSC. Reciprocal connections between the RSC, anterior thalamic nuclei (ATN), and the hippocampus constitute a limbic memory circuit that is essential for many forms of learning and memory.

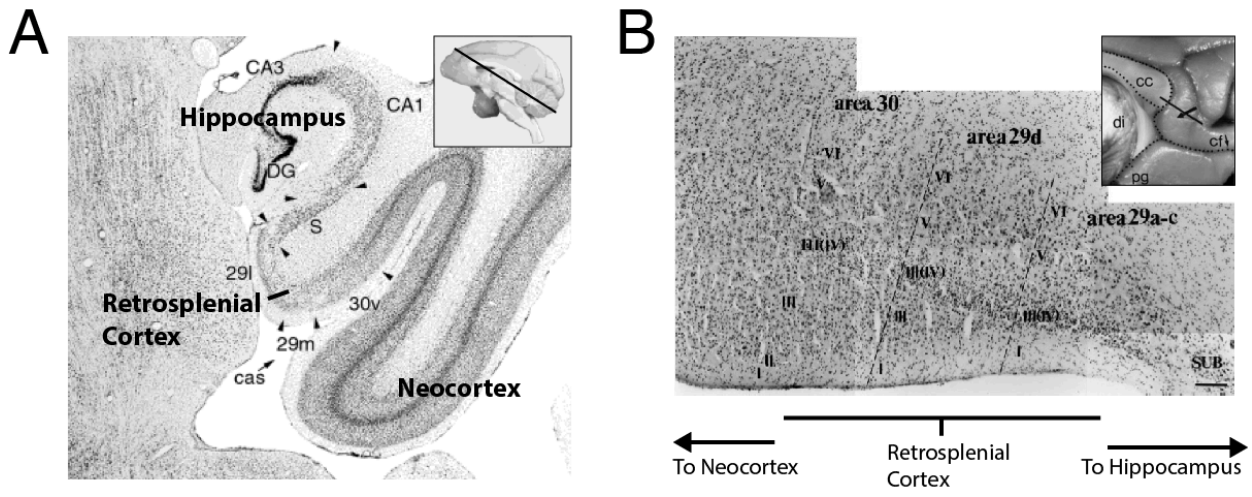

Figure 2 Cytoarchitectural transition between limbic and cortical regions via the retrosplenial cortex. The RSC (Brodmann areas 29 and 30) is often described as "transition" cortex due to its position between the 3-layer archicortex of the hippocampus and the 6-layer neocortex. Two examples from monkey brain slices. (**A**) The RSC divides an otherwise continuous strip of neural tissue running from the hippocampus to the neocortex. A low-magnification photomicrograph is shown of an oblique section through the RSC, 120 degrees from the coronal plane, of a macaque monkey. The position of the section is depicted in the inset. Abbreviations: S, subiculum; CA1, CA3, pyramidal fields of the hippocampus; cas, calcarine sulcus; DG dentate gyrus. Adapted from Koboyashi & Amaral (2000). (**B**) The transition from archicortex to neocortex is characterized by the division of layer III(IV), the granule cell layer, into layers III and IV, as well as the addition of a well-defined layer II. A light-field photomicrograph of a Nissle-stained section through the posteroventral aspect of the posterior cingulate gyrus of a rhesus monkey is shown. Section was taken at a level corresponding to the line in the inset. Scale bar indicates 100 micrometers. Abbreviations: SUB, subiculum; cc, corpus callosum; cf, calcarine fissure; di, diencephalon; pg, parahippocampal gyrus. Adapted from Morris, Petrides, & Pandya (1999).

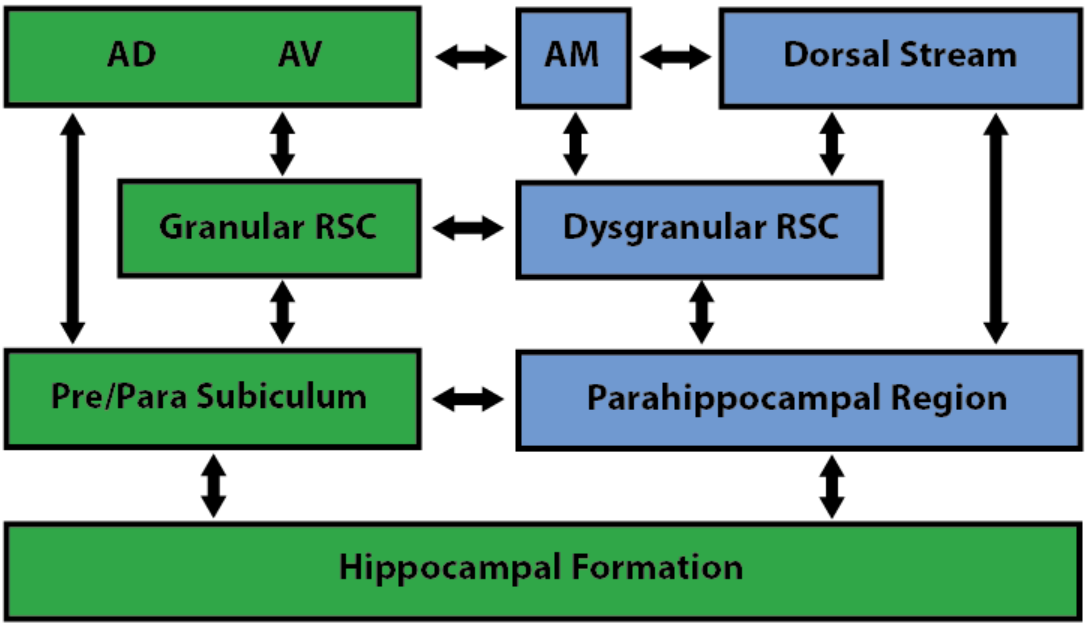

Figure 3 Schematic of RSC connectivity. The RSC shows regional differences in connectivity with limbic (green) and cortical (blue) regions. The granular RSC (areas 29a-c) shows greater connectivity with limbic regions such as the subicular cortex and the antero-dorsal (AD) and antero-ventral (AV) nuclei of the ATN, while the dysgranular RSC (area 30) shows greater connectivity with cortical regions including the parahippocampal region, the posterior parietal cortex, and early visual areas. The connections of the ATN also differ by region, with the AD and AV nuclei showing greater connectivity with limbic areas and the antero-medial (AM) nucleus showing greater connectivity with neocortical regions.

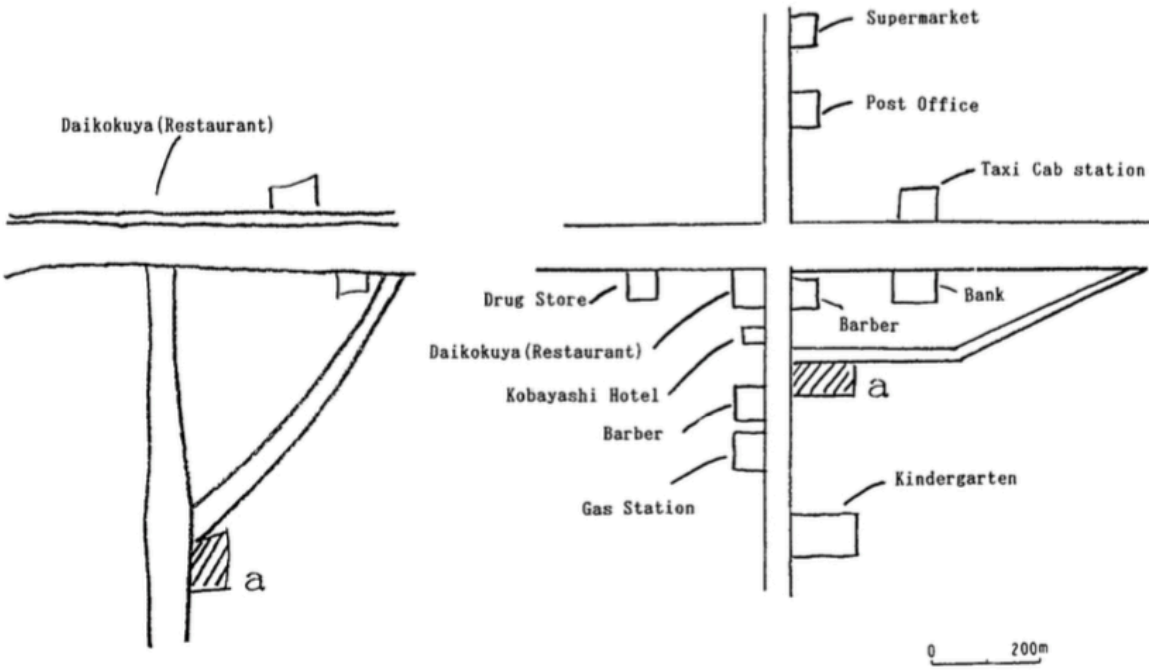

Figure 4 Impaired map drawing by patient with RSC lesion. Patients with RSC damage are severely impaired at drawing maps from memory. This patient (Patient 1, Takahashi et al., 1997) was incapable of recalling the locations of buildings in his neighborhood, despite being able to recall their names (left drawing). His wife's drawing is shown on the right for comparison. Adapted from Takahashi et al. (1997).

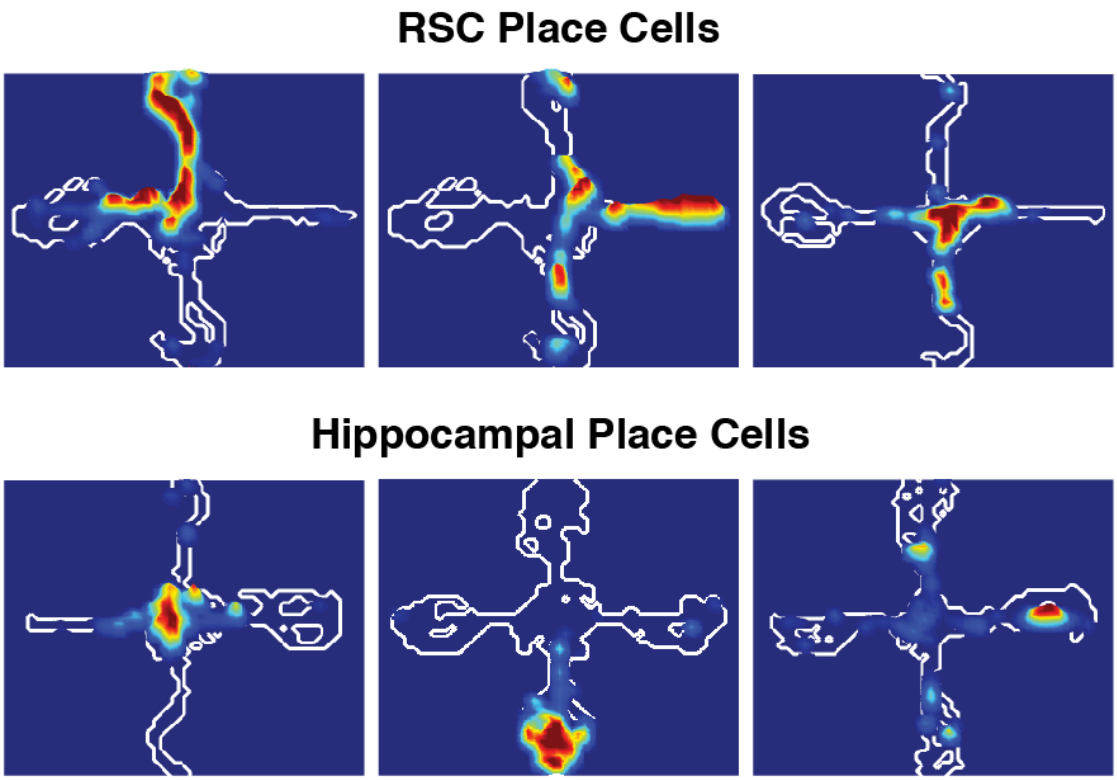

Figure 5 Spatial coding by RSC and hippocampal neurons during navigation. Three examples each of RSC and hippocampal place cells recorded during a blocked spatial alternation task on a plus maze. White lines indicate outer bounds of maze pixels visited by the rat during the recording session. Warmer colors indicate higher firing rates in that pixel. Adapted from Smith et al. (2012).

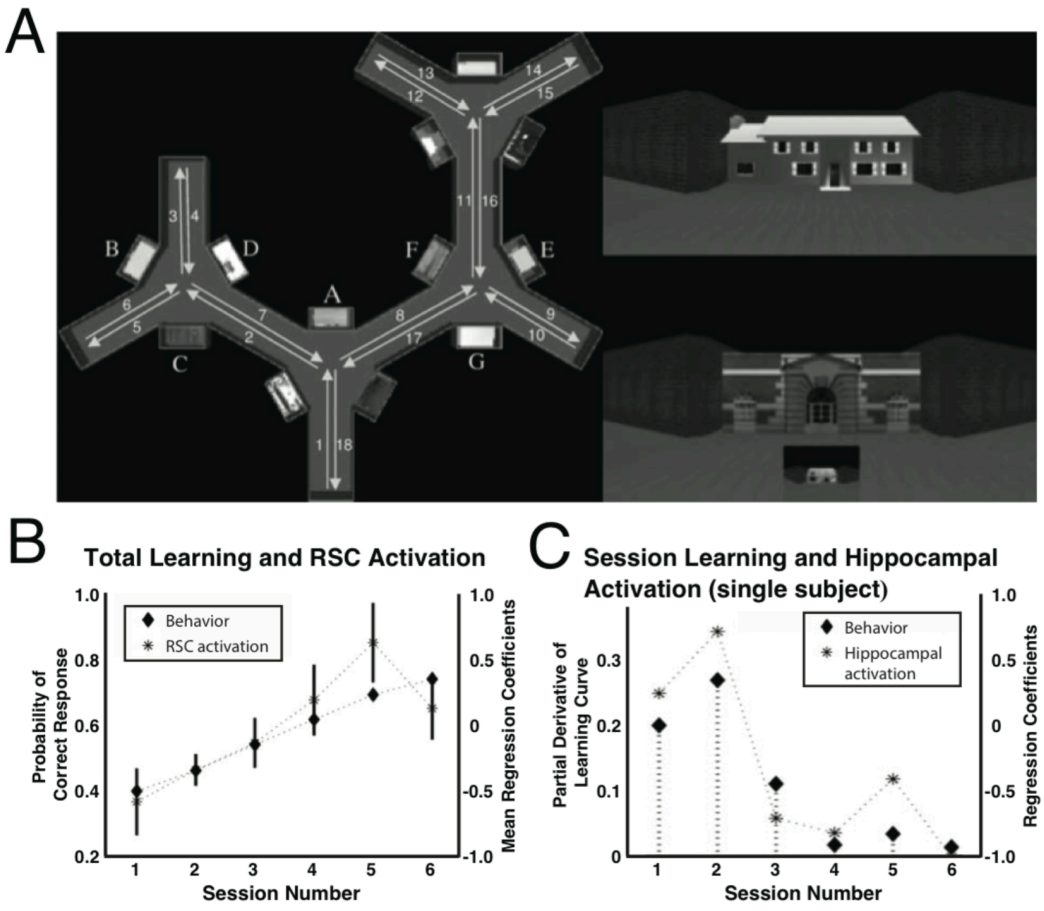

Figure 6 RSC and hippocampal activation during spatial navigation learning. (**A**) Participants received first-person tours of a virtual maze. On the left is a bird's-eye-view diagram of the maze (not shown to participants) showing the locations of the landmarks and the route of the tour. In the upper right is an example of the first-person view of a landmark. The lower right shows an example test question. Participants indicated by button press the relative position of the small building, imagining that they were standing in front of the large building. (**B**) RSC BOLD activation increased along with behavioral accuracy over the course of training. (**C**) In contrast, hippocampal activation was specifically correlated with the *slope* of the learning curve (data from one participant shown). Adapted from Wolbers and Buchel (2005).

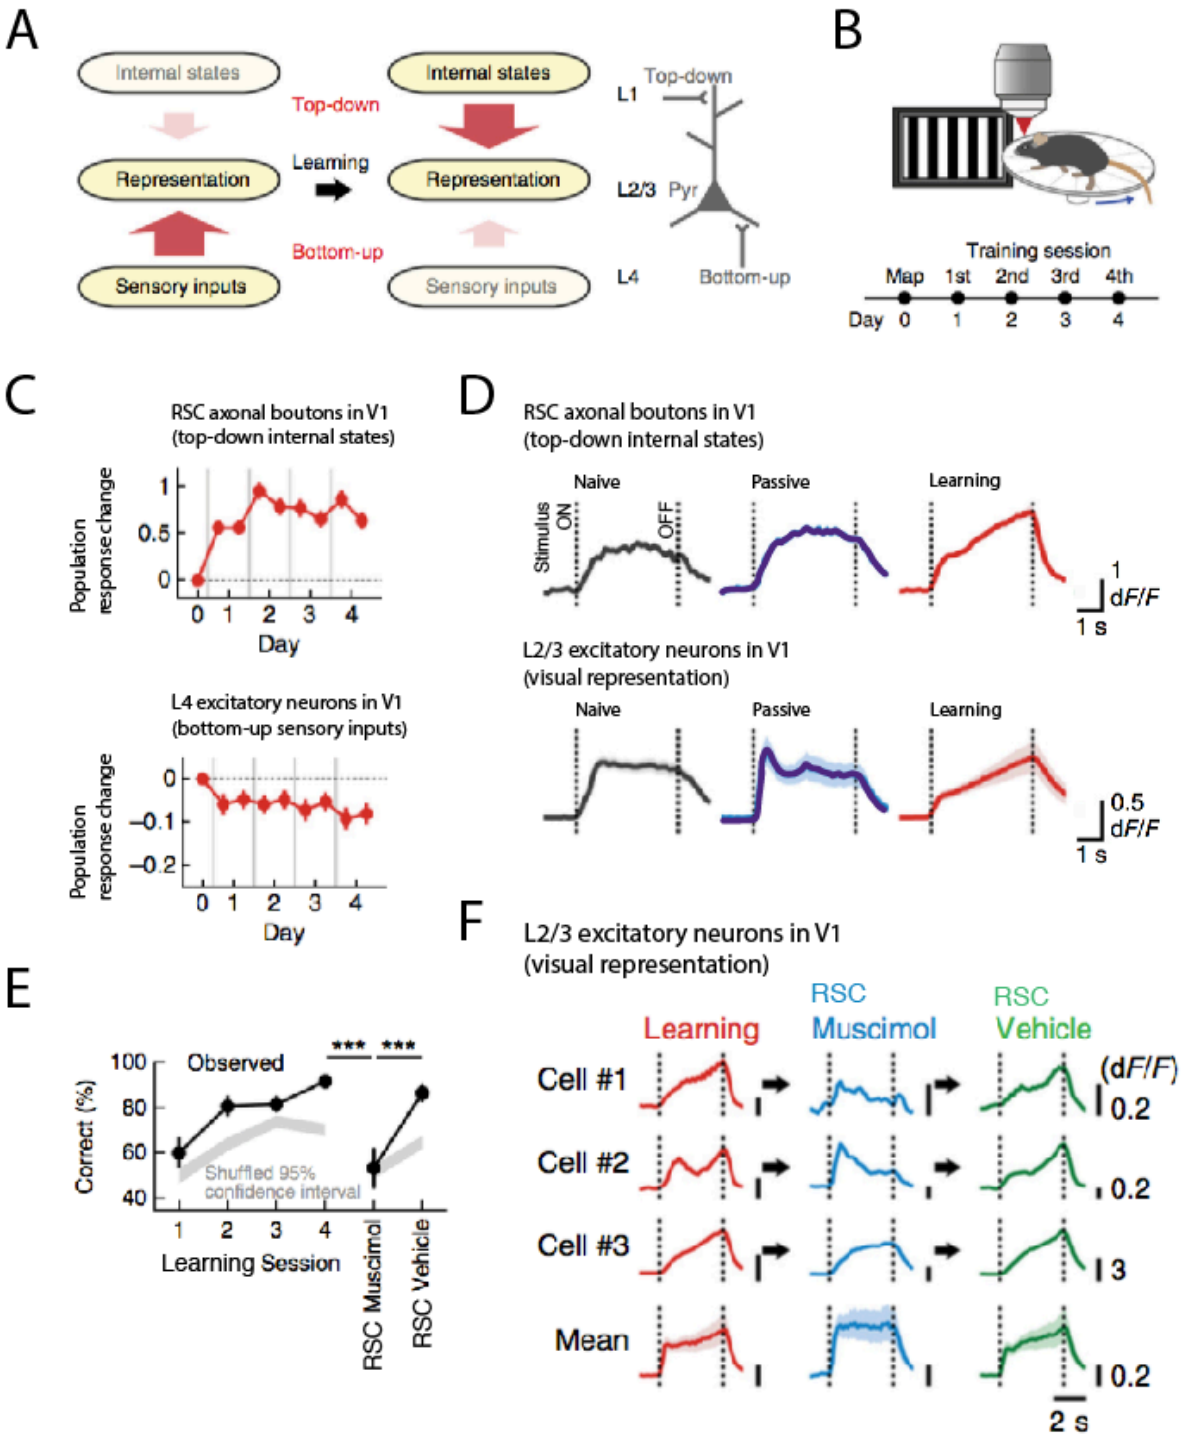

Figure 7 RSC neurons impose learned representations on V1. (**A**) A diagram showing the author's hypothesis that bottom-up inputs dominate in the naïve condition, while learning induces a top-down dominant states. (**B**) Top, schematic of behavioral set-up. The task required

head-fixed mice to run on a wheel if they detected the target visual grating stimulus in order to avoid a tail shock. Bottom, the training schedule. Recorded cells were identified on day 0 (also used for Naïve response baseline), followed by four days of training in one of 2 groups: learning, which learned the behavioral task, and passive, which saw the same visual stimuli but received no shocks. (**C**) Population response change of RSC axonal boutons terminating in L1 in V1 (top) and L4 excitatory neurons in V1 (bottom) over learning days. The value at each time point is the average population response (mean dF/F) during stimulus presentation on that day minus the response on day 0. dF/F refers to the change in florescence of individual neurons during an event, and is defined as the difference between event-related florescence and baseline florescence minus baseline florescence. (**D**) Mean of the normalized dF/F of all neurons responsive throughout the stimulus presentation period (dotted lines). Note the emergence of "ramp-up" firing patterns in RSC inputs and L2/3 excitatory neurons after learning. (**E**) RSC inactivation impaired task performance in well-trained mice. Performance improved over the four learning sessions. Muscimol infusions into the RSC reduced performance to chance-levels. Performance after muscimol infusions was significantly worse than after vehicle infusions. (**F**) Responses of three example L2/3 (representation) neurons and mean of all responsive L2/3 neurons showing that RSC inactivation eliminated the ramp-up firing patterns that had developed in L2/3 over the course of learning. Adapted from Makino & Komiyama (2015).

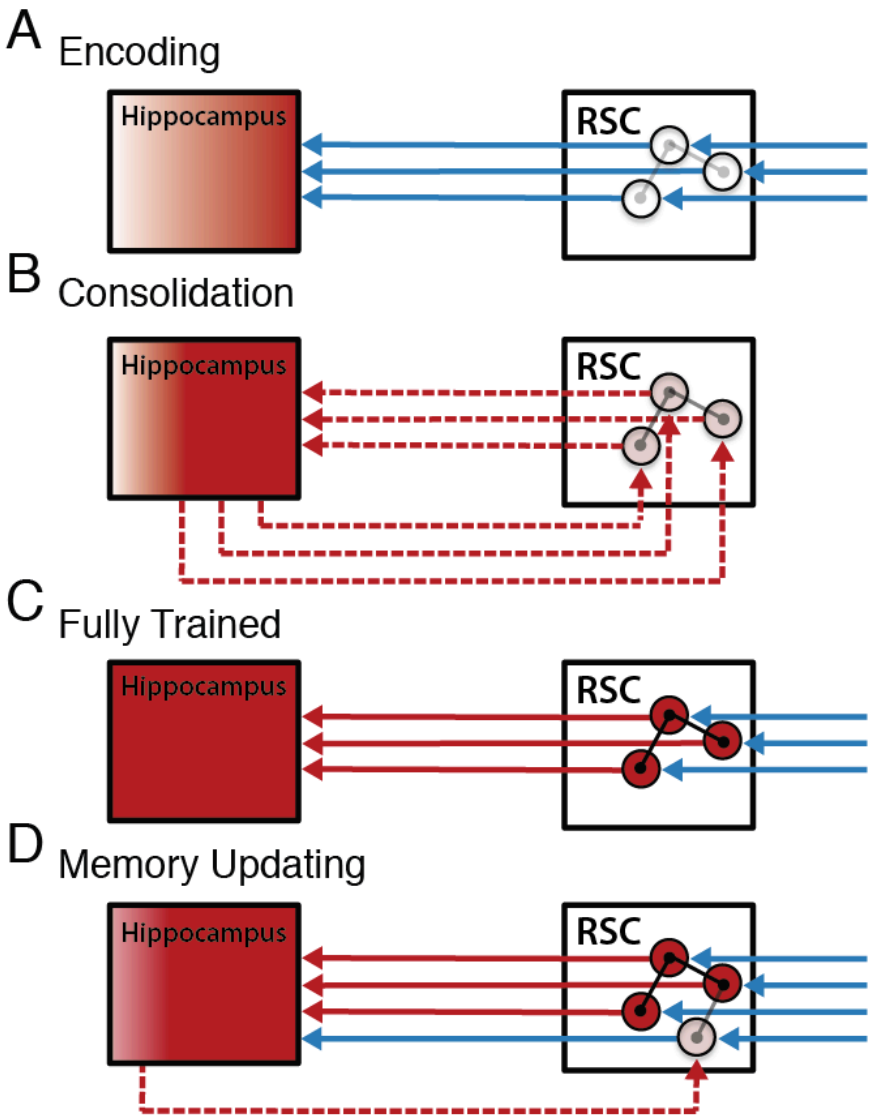

Figure 8 A model of RSC-hippocampal interactions over the course of learning. (**A**) Encoding in naïve animals is characterized by feedforward sensory information (blue lines) driving plasticity in both the cortex and the hippocampus. Plasticity in the RSC (circles and their connections) includes the early formation of cue-related activity, while plasticity in the hippocampus leads to the formation of a stable hippocampal memory representation (red). (**B**) After the learning event, interactions between the RSC and the hippocampus (dashed red lines) consolidate memories into a more stable form. (**C**) In the fully trained subject, feedforward sensory input activates the

consolidated RSC memory representation, which, in turn, activates the corresponding representation in the hippocampus. (**D**) Coordinated activity in the RSC and hippocampus enables memory updating, whereby the detection of novelty by the hippocampus initiates rapid consolidation of this new information into the existing cortical memory trace.

CHAPTER TWO

Context Encoding by Neural Ensembles in the Retrosplenial Cortex

Abstract: The RSC plays a central role storing context memories and processing contextual stimuli. However, next to nothing is known about how RSC neurons encode context or what kinds of stimuli are encoded in RSC context representations. We recorded RSC neuronal activity as rats investigated environmental contexts defined by distinct combinations of colors, sounds, and odors, as well contexts defined by combinations of object stimuli. The vast majority of individual RSC neurons showed both context-specific spatial representations and firing rate codes. Furthermore, RSC ensembles "drifted" steadily over the session, enabling us to decode current position in both space and time with high accuracy. Investigations of RSC ensemble responses to object contexts revealed a preference for spatial arrangements.

Introduction

Memory retrieval is often easier in the context in which the memory was acquired. For example, people often have an easier time recalling events from their youth when they are in their childhood home. Indeed, even thinking about the location where information was originally learned can improve retrieval (S. M. Smith, 1979). This context-dependence depends critically on the hippocampus, which plays an essential role in memory formation (Katche et al., 2010) and retrieval (Liu et al., 2012). However, some LTM remain context dependent even after they are consolidated out of the hippocampus (Cowansage et al., 2014), suggesting that context-dependence may be a more general property of memory throughout the brain.

Brain-wide investigations of context processing, such as studies using fMRI in humans (Bar & Aminoff, 2003) and studies of IEG activation in rodents (Wheeler et al., 2013), consistently point to the RSC as a critical node in the brain's memory and context processing networks. The RSC is intimately interconnected with the hippocampus (van Groen & Wyss, 2003), and is a site for the long-term storage of hippocampal dependent memories (Katche et al., 2013). Lesions of the RSC in humans impair memory for the layout of spatial environments (Takahashi, Kawamura, Shiota, Kasahata, & Hirayama, 1997), and lesions of the RSC in rodents impair both the formation and retrieval (Keene & Bucci, 2008a, 2008b; Robinson, Keene, Iaccarino, Duan, & Bucci, 2011) of spatial context fear memories. Recently, it was shown that optogenetic reactivation of the RSC can drive the retrieval of a context-dependent memory (Cowansage et al., 2014).

Despite this large body of research implicating the RSC in context processing, next to nothing is known about the nature of RSC context representations. In the hippocampus, where context representations have been extensively studied, individual neurons encode contexts

through their context-specific spatial firing patterns (i.e. place cells; for a review see D. M. Smith & Bulkin, 2014). However, it is unclear whether individual RSC neurons encode space at all (Alexander & Nitz, 2015, 2017; Cho & Sharp, 2001; Vedder, Miller, Harrison, & Smith, 2016), let alone in a context-specific manner. Nothing is known about how RSC ensembles, which appear to encode events and spatial locations with surprising accuracy (Alexander & Nitz, 2015, 2017; Vedder et al., 2016), may also encode contexts. Similarly, nothing is known about whether the RSC encodes temporal context (Bowers, Verfaellie, Valenstein, & Heilman, 1988; Todd, Meyer, & Bucci, 2015), as is seen in the hippocampus (Manns, Howard, & Eichenbaum, 2007) and medial prefrontal cortex (Hyman, Ma, Balaguer-Ballester, Durstewitz, & Seamans, 2012), or which environmental factors are used by the RSC to differentiate between two similar contexts. To answer these questions, we recorded RSC neuronal activity while rats explored a series of environmental contexts, and then compared the activity patterns of RSC ensembles between spatial locations, in distinct environments, and at different times throughout each session.

Results

RSC neural populations show context-specific spatial coding.

We recorded a total of 305 RSC neurons in four rats as they visited distinct spatial environments. Our recordings targeted the granular b subregion of the RSC (Rgb) bilaterally, although some neurons from the dysgranular RSC (RSA) were also included (see methods and **Figure 9** for anatomical details). On half of the recording days, rats made two visits to a black context and two visits to a white context in pseudo-random order. Contexts were 1m x 1m x 0.5 m boxes distinguished by distinct visual, auditory, and olfactory cues (see methods). Although a few neurons showed spatial firing localized to one area of the maze (e.g., one third of the maze, **Figure 10A top**), many more showed less specific firing around the entire maze (**Figure 10B bottom**). To determine whether individual neurons reliably expressed distinct spatial firing in the black and white contexts, we calculated pixel-by-pixel correlations between firing rate maps from each context. Remarkably, the vast majority (82.12%) of RSC neurons showed more similar spatial firing between two visits to the same context (e.g., black and black) than between two visits to different contexts (black and white; **Figure 10B**). RSC within-context correlation coefficients (mean = 0.46) were similar to those seen in the hippocampus (~0.5 in a similar task, from Law, Bulkin, & Smith, 2016). However, RSC between-context correlations (mean = 0.32) were far higher than those seen in the hippocampus, where between-context correlations are near zero.

These widespread yet small preferences for within-context spatial firing similarities (**Figure 10B inset**) suggest that each neuron may be contributing only a small amount of context-specific spatial information toward a more general population-level representation. To therefore examine how well we could use firing activity across the RSC population to decode

spatial position within each context, we combined neurons from all rats into a single population, and calculated ensemble firing rate vectors for every 250 ms time window during visits to each spatial bin. We then used a minimum distance classifier (see methods) to sort time windows into spatial bins based on firing rate activity. Importantly, we evaluated how accurately we could sort the time windows depending on whether the classifier was trained on firing rate maps from the same context or from the opposite context. This tests whether RSC spatial representations are specific to one context, or whether the same representation could be used to locate the rat in either context. We found that RSC activity from one context could be used to determine the rats' current spatial location better than expected by chance in either context (Kolmogorov-Smirnov test on distributions of spatial distance errors compared to control distribution obtained by shuffling spatial bin labels, both $p < 0.001$), consistent with the idea that the RSC uses prominent spatial features common to both contexts (e.g., spatial geometry of the boxes) to determine spatial position. However, we were nonetheless able to identify the correct spatial bin far more accurately when time windows were classified according to spatial firing activity from the same context (accurate classification within-context = 33.82%, between-context = 12.84%; Kolmogorov-Smirnov test on distributions of spatial distance errors, $p < 0.001$; **Figure 10C,D**). This is the first evidence for context-specific spatial firing in the RSC.

RSC neural populations encode distinct contexts.

In addition to distinguishing between spatial positions, many RSC neurons also displayed highly distinct firing rate codes for each context (**Figure 11A**), with 71.23% of neurons showing a significant firing rate difference between black and white context visits (as determined by a main effect of context in a repeated measures ANOVA, see methods). As with spatial

correlations described above, this property was not specific to a subset of RSC neurons, but instead existed to varying degrees throughout the population (**Figure 11B**). To therefore evaluate the distinctiveness of RSC rate codes for each context at the population level, we combined neurons from all rats into a single population, and computed ensemble firing rate vectors for every 250 ms time window during visits to each context. To visualize ensemble activity over all time windows, we performed principal components analysis on all firing rate vectors, including rate vectors from the three inter trial intervals (ITI), and plotted the first 2 principal components (**Figure 11C**). Consistent with our single-unit analysis, these plots suggested that RSC populations clearly distinguished between the black and white contexts. To quantify this, we computed the Euclidean distance from each time window to the mean activity of each context (black and white), and found that time windows were reliably more similar to the mean of the same context than of the opposite context ($t(11519) = 244.40$, $p < 0.001$, **Figure 11D,E**). To then determine whether RSC population activity could be used to discriminate between the two contexts, we employed a minimum distance classifier (as above, see methods) to sort individual time windows into either black or white contexts, and found that time-windows were accurately classified 98.56% of the time ($p < 0.001$, compared to a control distribution obtained by shuffling context labels).

Most neurons in the RSC showed a relationship between firing rate and running speed, and between firing rate and acceleration (**Figure 12**). However, this relationship was largely independent of the observed context-related firing (**Figure 13**). To nonetheless ensure that none of the reported differences in RSC ensemble firing rate activity could be attributed to small behavioral differences between context visits, we constructed a control data set by estimating and correcting for the relationships between firing rate, velocity, and acceleration. In no case were

the reported outcomes explained by controlling for the influence of these variables, and, in most cases, the reported effects were strengthened by the controls.

RSC neural populations encode visit order.

Temporal features of experiences are essential for forming accurate memories of events, and recent work has suggested that temporal coding may be an important feature of the hippocampal memory system. To test whether RSC ensembles encoded temporal features of the session, we again combined neurons from all rats into a single population and assembled firing rate vectors for 250 ms time bins throughout the session. However, instead of matching time windows across sessions according to context type, we matched them according to which context visit (visits 1-4) the rat was in. Therefore, every time window contained roughly equal numbers of cells from visits to the black context and visits to the white context. To determine whether firing activity of RSC ensembles changed with time over the session, we calculated the Euclidean distance from the time window 2 minutes before the first context visit (while the rats were in the ITI cylinder) to every time window during each context visit. This revealed that distances steadily increased over the session as RSC ensemble activity drifted away from its starting state ($r = 0.52$, $p < 0.001$; **Figure 14A**). To therefore test whether RSC ensembles reliably distinguished between the four visits, we calculated the Euclidean distance between two visits to the same context and visits to the opposite context as a function of the lag between their visit positions (e.g., the distance between two visits to the white context that were 1, 2, or 3 visits apart; **Figure 14B**). A two way ANOVA (within/between contexts comparisons X three lag distances) revealed a main effect of context comparison ($F(1, 34655) = 1979.35$, $p < 0.001$), consistent with the context discrimination analysis above. Importantly, there was also a main

effect of lag distance ($F(2, 34655) = 1850.21$, $p < 0.001$), showing that RSC ensembles uniquely encode the temporal order of context visits. Additionally, a context comparison by lag distance interaction ($F(2, 34655) = 331.08$, $p < 0.001$) suggests that within-context differences are more stable over visits than are between-context differences.

The distances plotted in **Figure 14A** additionally suggested that the RSC ensemble may undergo a large shift between visits to the context apparatus and the ITI cylinder. Similar ensemble state shifts have been observed in the hippocampus (Jezek, Henriksen, Treves, Moser, & Moser, 2011), and are believed to enable a rapid change in the representation of two different environments. To quantify this, we calculated the rate of change in the ensemble state (i.e. drift velocity) as rats entered contexts from the ITI cylinder, and found that maximum drift velocity occurred immediately before, and reached its lower limit immediately after, the rat was placed in the context (**Figure 14C top**). Similarly, the distance between the instantaneous RSC ensemble activity and the mean of the to-be-visited context declined in the seconds before the context visit, and reached its lower limit shortly after (**Figure 14C bottom**). Plots of the first two principal components of ensemble activity over the session similarly revealed a large shift in the stage of the RSC ensemble between context and ITI visits followed by a relatively stable state within that visit (**Figure 14D**).

Together, these data suggest that the RSC may be encoding the temporal order of context visits in addition to the position of the rat within specific contexts. To test, then, whether we could use instantaneous activity in the RSC ensemble to decode the spatial position of the rat within both the context (i.e. environment) and the particular visit (i.e. ordinal time), we combined neurons from all rats into a single population, and calculated ensemble firing rate vectors for every 250 ms time window during visits to each spatial bin in each of the four spatio-

temporal context visits (black visit 1, black visit 2, white visit 1, white visit 2). We then computed the Euclidean distance from each time window to the mean activity of every spatial bin in each spatio-temporal context and sorted time windows into the most similar spatial bin based on firing rate activity (i.e. minimum distance classification, see methods). Remarkably, RSC activity could be used to determine the rats' current spatio-temporal spatial location 32.22% of the time (chance = 0.69%, binomial test, $p < 0.001$; **Figure 14F**). When classification errors did occur, they generally occurred within the correct spatio-temporal context, such that we could identify the correct spatio-temporal context 97.59% of the time. Together, these data show that the RSC distinctly encodes both the spatial and temporal elements of an experience.

RSC neural populations preferentially encode spatial arrangements.

The RSC plays an important role encoding landmarks (Auger, Mullally, & Maguire, 2012) and other navigation cues (Vedder et al., 2016), but it is unknown how this function relates to the RSC role in context processing. To investigate this, we recorded in the same group of rats as they explored environments that differed in the arrangement or the appearance of landmark objects (painted 10" tall terra cotta cones). The environments included two object-arrangement conditions, where the same objects were arranged differently in each condition, and two object-appearance conditions, where two distinct sets of objects shared the same arrangement (**Figure 15**). We found that individual neurons developed highly specific firing rate codes for each of the four object conditions (**Figure 16A**), with individual neurons showing varying degrees of selectivity for both the arrangement and appearance manipulations. This suggests that RSC ensembles uniquely encoded these conditions similar to the encoding of black and white contexts described above. To visualize RSC population-level coding of the four environments, we

combined neurons from all rats into a single population, and computed ensemble firing rate vectors for every 250 ms time window during visits to each of the four object conditions, and then plotted the first 3 principal components (**Figure 16B**). These plots suggested that RSC populations distinctly represented each of the two arrangement conditions, while distinguishing between the two appearance conditions only minimally. To quantify this, we computed the Euclidean distance from each time window to the mean activity of each of the four object conditions, and found that RSC ensemble activity was more distinct between arrangement conditions than between appearance conditions ($t(2879) = 8.27$, $p < 0.001$, **Figure 16C,D**). To then determine whether RSC population activity could be used to discriminate between the four object conditions, we employed a minimum distance classifier (as above, see methods) to sort individual time windows into object conditions, and found that the time-windows were accurately classified 95.26% of the time (binomial test, $p < 0.001$). Consistent with the idea that the RSC ensemble representations of appearance conditions are more similar than the representations of arrangement conditions, time windows were misclassified between the two appearance conditions more than twice as often as between the arrangement conditions (classification errors within-arrangement = 207, classification errors within-arrangement = 85). The representational similarity between the two appearance conditions was not due to rats' inability to distinguish between them, as a separate group of rats similarly distinguished between object-appearance and object-arrangement manipulations in a novel context detection task (**Figure 17**).

Context representations in granular and dysgranular RSC.

Very little is known about how RSC sub regions, such as the granular RSC (Rgb) and dysgranular RSC (RSA), may differentially contribute to spatial and contextual representations. However, anatomical data suggest that Rgb might play a greater role due to its stronger connectivity with the hippocampus (van Groen & Wyss, 1992, 2003). While the majority of neurons in our data set were recorded in Rgb (245), a sizeable minority (60) was recorded in the RSA. We therefore compared these regions on each of our main research questions: spatial coding of distinct environments, rate coding of distinct context, temporal coding of visit order, and selectivity for spatial arrangements. To first test whether there were differences in spatial coding of different environments, we performed a two-way ANOVA with region (Rgb and RSA) and context comparison (between and within) as factors, and found significant main effects for both region ($F(1, 286) = 5.97$, $p < 0.05$) and context comparison ($F(1, 286) = 16.19$, $p < 0.001$), but no context by region interaction ($F(1, 286) = 1.60$, $p = 0.21$), suggesting that the two regions show similarly unique spatial codes between the two contexts (**Figure 18A**). Post hoc comparisons confirmed that both the Rgb ($t(111) = 8.34$, $p < 0.001$) and the RSA ($t(32) = 3.37$, $p < 0.01$) showed significantly greater correlations within than between contexts.

To then test the populations measures of context coding, temporal order coding, and selectivity for spatial arrangements, we constructed separate ensemble firing rate vectors from neurons recorded in Rgb or RSA, and then calculated the distance between contexts, visits, or object conditions as in Figures 2-4. To test for regional differences in rate coding of black and white contexts, we performed a two-way ANOVA with region (Rgb and RSA) and context comparison (between and within) as factors, and found significant main effects for region ($F(1, 46076) = 205.06$, $p < 0.001$) and context comparison ($F(1, 46076) = 1129.91$, $p < 0.001$), as well as a region by context interaction ($F(1, 46076) = 58.53$, $p < 0.001$; **Figure 18B**). This shows that

while both regions differentiated between context visits, the Rgb differentiated more than did the RSA, especially between the black and white contexts.

To test for regional differences in rate coding of temporal order, we performed a two-way ANOVA on the regions (Rgb x RSA) and time between context visits (lags 1-3), and found significant main effects for both region ($F(1, 69322) = 103.81$, $p < 0.001$) and time ($F(2, 69322) = 760.84$, $p < 0.001$), as well as a region by lag interaction ($F(2, 69322) = 71.14$, $p < 0.001$; **Figure 18C**). This showed that while both regions differentiated between the three ordered context visits, the Rgb differentiated more than did the RSA, especially between visits at the beginning and end of the session (a lag distance of 3).

Lastly, to test for regional differences in rate coding of object arrangement and appearance, we performed a two-way ANOVA with region (Rgb and RSA) and object condition (arrangement and appearance) as factors, and found a main effects of region ($F(1, 11516) = 54.11$, $p < 0.001$), and a region by context interaction ($F(1, 11516) = 39.68$, $p < 0.001$; **Figure 18D**). Post hoc comparisons revealed that while the Rgb showed greater differentiation between object arrangements than between object appearances ($t(2879) = 6.73$, $p < 0.001$), the RSA showed the opposite relationship (appearance > arrangement, $t(2879) = 4.03$, $p < 0.001$).

Discussion

The RSC plays an important role in spatial cognition and is heavily interconnected with the hippocampus and other regions that process memories in a context-dependent manner. Human fMRI work has pointed to the RSC as the center of a network for context processing, and recent work in rodents has identified the RSC as a major contributor to the long-term storage of context memories. However, no research has directly examined the form of RSC context representations, and it is unknown how RSC neural processing might supports these many functions. Our findings bridge this gap by providing the first clear evidence that RSC ensembles encode context-specific information, including distinct spatial representations and ensemble rate-codes for both environmental and temporal contexts. Additionally, we found that RSC ensembles were more sensitive to manipulations of the spatial layout of the environment than to its visual features, providing new insight into the recently identified role for the RSC in landmark processing.

Relative to the hippocampus, little work has focused on RSC context representation. In the hippocampus, context representations take the form of context-specific spatial firing among hippocampal neurons (i.e. place cells), yet recent studies of RSC spatial coding have consistently failed to identify hippocampal-like place cells in the RSC (Alexander & Nitz, 2015, 2017; Vedder et al., 2016). The present data were largely consistent with these recent studies in that individual RSC neurons did not show clearly isolated firing fields. Nonetheless, we found that most (82.12%) RSC neurons showed at least some context-specific spatial firing, and that, at the population level, RSC neuronal ensembles showed far superior spatial decoding within a context than between two different contexts. These are the strongest data yet showing that RSC neurons

encode spatial position, and are the very first data showing context-specific spatial mapping (similar to 'remapping' in the hippocampus).

In addition to this spatial coding, RSC neurons were remarkably sensitive to which environmental or temporal context the rat was currently occupying. We identified that ~70% of RSC neurons showed significantly different firing rates for each context, and we were able to use ensemble activity to identify the rat's current context with near-perfect (98.56%) accuracy. We similarly found that RSC ensembles "drifted" over the course of the session, such that the instantaneous firing rate could also be used to decode the rat's current temporal position in the order of context presentations (similar to what has been seen in the prefrontal cortex, Hyman et al., 2012). The spatial, contextual, and temporal codes were largely independent of one another such that we could identify the rats' position in both time and space with remarkable accuracy. Very little work has considered the role of the RSC in temporal processing (Robinson et al., 2011; Todd et al., 2015). However, in the hippocampal system, the combination of spatial and temporal firing codes is believed to provide the substrate for encoding and retrieval of memory episodes (Eichenbaum, 2014). A similar function for the RSC would be consistent with both its prominent position in the brain's default memory network (Spreng, Mar, & Kim, 2009) as well as with its interconnectivity with the hippocampus (van Groen & Wyss, 1990, 1992, 2003).

Lastly, we investigated the spatial factors driving RSC context representations by evaluating the influence of RSC sensitivity to contextual objects (Bar & Aminoff, 2003) and landmarks (Auger et al., 2012; Auger, Zeidman, & Maguire, 2015). We utilized the RSC ensemble context coding described above to show that the RSC is far more sensitive to spatial arrangements of objects than to specific objects per se. This is consistent with its role in spatial representation and with its anatomical connections with the dorsal cortical processing stream

(Miller, Vedder, Law, & Smith, 2014), which preferentially encodes spatial stimuli. Although this area of research remains at an early stage, these findings suggest that the RSC role in landmark processing may simply be another facet of a broader role encoding contexts.

Methods

Subjects and Surgery

Subjects were 4 adult male Long Evans rats (Charles River Laboratories, Wilmington, MA) weighing 250g-300g upon arrival. Rats were placed on a 12/hr/12hr light/dark cycle with lights on at 8am and were allowed to acclimate to the vivarium for at least one week prior to surgery. Rats were implanted with a custom-built electrode microdrive (Macdonald et al. 2011) under isoflurane anesthesia. The microdrive contained 20 moveable tetrodes made from twisting four 17μm platinum/iridium (90%/10%) wires, platinum plated to an impedance of 50-300kΩ, and arranged in two (one in each hemisphere) 10-tetrode linear arrays angled 30 degrees toward the midline and spanning approximately 4mm along the rostrocaudal axis of the brain (2-6mm posterior to Bregma, ±1.5mm lateral). In two subjects, an additional stainless steel local field potential probe was implanted into the hippocampal fissure (a 200μm probe in the left hemisphere of subject R1917 at -3mm posterior to Bregma, and a 50μm probe in the right hemisphere of subject R1920 at -4.5mm posterior to Bregma, both 2.5mm lateral). All rats were given 7 days to recover from surgery, and then placed on food restriction until they reached 80-85% of their free-feeding weight. Water was always available ad libitum. Over the course of several days, Tetrodes were lowered into the RSC (35-70μm daily) until a depth of at least 1 mm was reached. All tetrodes were lowered after each recording session (~15-30μm) in order to maximize the number of unique neurons recorded.

Apparatus

Rats foraged for sprinkles in two distinct contexts (PVS boxes measuring 1m X 1m X 0.5m deep) that differed in the color of the box (black or white), the color of the surrounding

walls/curtains (white or black), background masking noise (pink or white noise), and odor (before training and during every intertrial interval, boxes were wiped with either scented or unscented baby wipes, Rite Aid, Inc). On approximately half of the recording days, subjects visited the same environmental context (black or white) for four consecutive visits, and objects (painted terra-cotta cones, 10" tall, 6" diameter) were placed on the floor of the box to create four unique environments (Supplementary Figure 3A-B) each day. The unique environments were designed to evaluate the contributions of object appearance (object A, B, or C) and position (arrangement 1, 2, or 3) to RSC context representations. The specific environments were (1) object A in arrangement 1, (2) object A in arrangement 2, (3) object B in arrangement 3, and (4) object C in arrangement 3. The environmental context was alternated across days to ensure that an equal number of object sessions occurred in each context.

Acclimation Procedure

Before surgery, rats were acclimated to each of the environmental contexts and object manipulations. Acclimation consisted of 2 days of visits to the two environmental context conditions presented in ABAB order (counterbalanced across days), and 12 days of visits to the object contexts (6 in the white box and 6 in the black box) presented in pseudo-random order. During acclimation sessions, rats were given four 12 min context visits with five ~3 min ITI periods (before, between each visit, and after). During ITI periods, the rat was placed in an opaque cylinder (30 cm height, 65 cm height) lined with bedding material and closed at the top except for a small (5 cm diameter) hole. Transitions between the cylinder and the box (or the reverse) began by turning on (or off) the background music, opening the ITI cylinder, lifting the

rat and placing him into the center of the box (or back into the ITI cylinder). Transitions lasted ~10 seconds.

Recording Procedure

Prior to each recording session, rats were plugged into the recording equipment and allowed to acclimate by exploring a circular environment (diameter 1m) not used during the experimental procedures for ~10min. Recording began when the rats were placed into the ITI cylinder. Recording procedures were the same as acclimation procedures described above. Digital flags were manually entered immediately after the rat was placed into the box to indicate the beginning of the context visit, and immediately before the background music was turned off to indicate the end. ITI periods were defined as the periods between context visits.

Recordings

Neuronal spike data and video data were collected throughout learning (Digital Cheetah Data Acquisition System, Neuralynx, Inc. Bozeman, MT) and stored to disc along with timestamps for offline sorting (SpikeSort3D, Nueralynx, Inc.). The rat's position and head direction was monitored by digitized video of green and red tape attached to the rat's headgear.

Histology

After completion of the recording experiment, rats were transcardially perfused with 4% paraformaldehyde in phosphate buffered saline. Brains were removed and stored for at least 24hrs in 4% paraformaldehyde before being transferred to a 30% sucrose solution for storage until slicing. Coronal sections (40μm) were stained with 0.5% cresyl violet to locate tetrode

tracks. Tetrode positions during recording sessions were identified using depth records noted during tetrode lowering and tracks observed in the stained tissue (Supplementary Figure 1). Boundaries of RSC regions were determined in accordance with The Rat Brain in Stereotaxic Coordinates (Paxinos and Watson, 2013). Neuronal records from tetrodes located outside of the RSC were excluded from the data set.

Neuron counts

A total of 305 RSC neurons were recorded from four rats. Roughly half (146) of the neurons were recorded on days when the rats visited black and white contexts, while the rest (159) were recorded during object condition days.

Data Analyses

Spatial Correlations. Spatial firing rate maps for each neuron were constructed by summing the total number of spikes that occurred in a spatial bin (2.5 x 2.5 cm) and dividing by the amount of time spent in that bin. Bins visited on only one trial were excluded. Rate maps were smoothed using an adaptive method designed to optimize the trade-off between blurring error and sampling error (Skagg, Wilson, McNaughton, & Barnes, 1996). Pixels that were not visited during a context visit (2.92% +/- 2.36% SD of pixels during each session) were excluded from the smoothing procedure. Firing rates for these pixels were then estimated based on the firing of neighboring pixels via linear interpolation or extrapolation.

Single-unit pixel-by-pixel correlations of spatial firing between two context visits were calculated by vectorizing the context firing rate maps and then computing the correlation between the two context firing rate vectors. For each neuron, we computed both the mean

within-context correlation and the mean between-context correlation. The within-context correlation was the average of the correlation between the first and second visit to the white context (W1 & W2) and the correlation between the first and second visit to the black context (B1 & B2). The between-context comparison was the average of the four correlation coefficients resulting from pair wise comparisons between black and white context visits (B1 & W1, B1 & W2, B2 & W1, B2 & W2). The distribution of within-context correlations from every neuron was then compared to the distribution of between-context correlations using a paired t-test.

Minimum Distance Classification of Spatial Position. Ensemble firing rate vectors for every 250 ms time window during the recording session were constructed by combining neurons across subjects and sessions. Firing rates from every neuron were z-scored across all time windows to control for differences in baseline firing rates, which render variance and distance measures overly sensitive to the coding properties of relatively few neurons. Time windows from each session were then sorted according to which context (black and white) and spatial bin (6 X 6, approximately 17cm X 17cm) the rat was in during that time window. Because rats differed in number of visits to each pixel, only 40 time windows from each spatial bin in each context (20 from each visit, selected evenly from across the session) were included in the analysis. This resulted in 2880 time window firing rate vectors during each of the two contexts (6 pixels X 6 pixels X 40 time windows).

To then classify time windows into spatial bins, we employed a minimum distance classifier, whereby we first calculated the distance from each time window to the mean of each spatial bin (the mean of all time windows drawn from that spatial bin), and then the time window was classified according to the smallest distance (i.e. the time window was assigned to the spatial

bin with mean firing rate activity that was most similar to the activity occurring during that time window). For within-context classification, we found the distances to spatial bins occurring within the same context (excluding that time window; e.g., the distance from a time window from a white context visit to the mean of all white context spatial bins). For between-context classification, we found the distances to spatial bins occurring within the opposite context (e.g., the distance from a time window from a white context visit to the mean of all black context spatial bins). Classification errors occurred when the minimum distance was to a spatial bin other than the category from which the time-window was drawn. The spatial distance from the classified bin to the actual bin was the distance between bin centers.

Context and Visit Selectivity of Single Neurons. For each neuron, we calculated the firing rate from each 30 sec time window during each of the four context visits, and then performed a 2 (black and white contexts) X 2 (first and second visits) repeated measures ANOVA. Context selective neurons were defined as having a significant ($p < 0.05$) main effect of context. Visit selective neurons were defined as having a significant main effect of visit.

Population Distance Between Contexts. Ensemble firing rate vectors for every 250 ms time window during the recording session were constructed by combining neurons across subjects and sessions. Time windows from each session were sorted according to which context or ITI visit (ITI1, ITI2, ITI3) the rat was in during that time window. Because sessions had minor variations in the lengths of context visits and ITIs, only the first 12min from each context visit and the first 2 min from each ITI were included. This resulted in 2880 firing rate vectors from time windows during each of the four context visits (12 min X 240 time windows) plus 480 firing rate vectors

from time windows during each of the three ITI visits (2 min X 240 time windows). Firing rates from every neuron were z-scored across all time windows.

To calculate distances between contexts, we employed a minimum distance classifier as above. To do this, we used an iterative procedure whereby we excluded one time window from the data set, calculated the mean firing rate vector for each context from the remaining time windows, and then computed the Euclidean distance between the excluded time window and each context (similar results were obtained if the within distance only included time windows during the other visit to that context). Because the firing rates are z-scored, distances are displayed as +/- 1 z-scored unit, which is the standard deviation of firing rates from all time-windows. Under the null hypothesis (context visits are arbitrary groupings of time windows drawn from a single distribution), the average distance between any time-window to any context mean is 1 z-scored unit. Distances greater than 1 are larger than (i.e. the representations are less similar than) expected under the null and distances less than 1 are smaller than (i.e. the representations are more similar than) expected under the null.

Minimum Distance Classification. Classification of individual time windows into categories (e.g., the four object conditions) was based on the distance measurements calculated as above. For each time window, we first calculated the distance from that time window to the mean of each of the category means, and then the time window was classified according to the smallest distance (i.e. the time window was assigned to the category with mean firing rate activity that was most similar to the activity occurring during that time window). Classification errors occurred when the minimum distance was to a category other than the category from which the time-window was drawn.

Calculating Population Drift. To calculate population drift over the session, we computed the distance from the average ensemble activity occurring during the first minute of the first context visit to every other time window in the session (Figure 3A), and found the correlation between distance and time. Ensemble firing rate vectors for every 1s time window during the recording session were constructed as above by combining neurons across subjects and sessions. However, this analysis disregarded the specific context from which each time window was recorded, and therefore preserved only the order of each context visit (1-4). Population drift over the course of the session was then calculated by finding the correlation between distance and time over the session (0-58 min). Drift velocity at each time window as rats entered a context was calculated by finding the distance from the previous to the upcoming time window divided by time (2 s). Velocity was calculated from 30 s before to 120 s after the moment the rat was placed onto the floor of the context of each context visit, as manually flagged by the experimenter. The velocities were then averaged across the four visits. Distances from the current time window to the mean of the current context visit were also averaged across all four visits. All population drift analyses included both context and object sessions, but similar results were obtained if object conditions were excluded.

Comparisons Between Context Visits by Lag Distance. To measure the population distance between two context visits as a function of lag in visit position, we constructed firing rate vectors independently for each lag distance comparison (lags of 1, 2, or 3 positions), and then computed the standardized Euclidean distance between firing rate vectors from the first visit to firing rate vectors from the second visit. To construct firing rate vectors for each lag comparison, we

combined neurons across subjects and sessions that included the relevant lag comparison. For example, a session with the context order of black, white, white, black, would only contribute to within-context comparisons of visits that differed by 3 positions (black-black) or by 1 position (white-white). Firing rate vectors were created for every time window in each of the two visits, and then distances between the two context visits were computed at every time window by calculating the distance between that time window and the mean of the other visit. Spatial correlations between context visits as a function of lag distance were calculated in a similar manner, but instead of calculating distances between rate vectors at every time window, we found the correlation between rate vectors at every spatial bin (see *Spatial Correlations* above). These analyses only included black and white context sessions.

Spatial-temporal Classification. Ensemble firing rate vectors for every 250 ms time window during the recording session were constructed by combining neurons across subjects and sessions (only including the first 12min from each context visit). Time windows from each session were then sorted according to which context and visit (black visit 1, black visit 2, white visit 1, and white visit 2) the rat was in during that time window. To determine spatial firing, we overlaid a 5 X 5 spatial grid over each context and found the first 40 time windows from each context visit that occurred while the rat was occupying each spatial bin. We selected small (250 ms) time windows to reduce the possibility that the rat occupied more than one spatial bin during a single time window, and we selected 25 spatial bins and 40 time windows in order to balance the tradeoff between having large numbers of time samples with having to exclude sessions where the rat spent insufficient time in each bin. This left us with 4000 time windows (40 time windows X 25 spatial bins X 4 contextual-temporal bins), each arising from one of 100 spatial-temporal

bins (25 spatial bins X 4 contextual-temporal bins). Time windows were then classified into one of the spatial-temporal bins using a minimum distance classifier as described above. For each time window, we calculated the spatial distance between the actual location of the rat and the location output of the classifier (irrespective of context) by calculating the spatial Euclidean distance between bin centers within the apparatus. This analysis only included black and white context sessions because object conditions did not include repeated visits to the same context.

Object Conditions. Distances between object and arrangement conditions (Fig. 1d) were computed similarly to comparisons between black and white contexts. Time windows from each object appearance condition were used to calculate one appearance distance, and time windows from arrangement conditions were used to calculate one arrangement distance. Minimum distance classification was conducted as above.

Novel Context Detection. Eight un-operated rats not used in the electrophysiology study were tested for their ability to detect the different object manipulations using a variant of the novel object recognition task. Each rat visited four contexts a day for two days, with each day testing the ability to recognize a change in either object arrangement or object appearance. In the object arrangement test, the rat visited the same arrangement manipulation (e.g., object A in arrangement 1) over three consecutive context visits, and then saw the opposite arrangement manipulation on the fourth visit (e.g., object A in arrangement 2). The same procedure was used for the object appearance tests, except with objects B and C in arrangement 3. All context visits were recorded to video, and the number of object explorations was later counted and used as a measure of interest in the objects. Object explorations were defined as direct contacts with the

mouth, nose, or paw, lasting at least one second, and did not include contacts judged to be accidental, such as backing into the object. In addition, contact in which the rat was leaning on the object as a way of exploring other aspects of the spatial context was not counted as exploration. Only the first 3 min of each visit were included in the analysis to avoid including time periods after the rat had become familiar with the novel object.

Velocity and Acceleration Control. We constructed a control data set by estimating and correcting for the relationships between firing rate and velocity and between firing rate and acceleration. For each session and cell, we calculated the firing rate and velocity of the rat during every 1 s time window during context visits, and then calculated the mean firing rate in each of 30 equally sized velocity bins spanning the range of observed velocities. We estimated the relationship between firing rate and velocity by fitting a quadratic function to the mean firing rates, and then used the function to determine the correction required to increase or decrease firing rates in every bin in order to render firing rates, on average, equal to those seen at 0.4 m/s. This velocity was chosen because the relationship between firing rate and velocity appeared to approach an asymptote around 0.4m/s for most cells. We then applied these corrections to every time window by deleting the original spike events occurring during that time window and replacing them with a number of linearly spaced spike events equal to the original number of events plus the correction. In order to create the control data set, we repeated this procedure for every cell in every session from every subject in the original data set. The entire procedure was then performed a second time on the new control data set to correct for the remaining relationship between firing rate and acceleration (corrections were to 0.4 m/s^2). All analyses

presented in this paper were repeated on the both the original and this control data set, and similar results were obtained in every case.

Comparisons between RSC Subregions. The original data sets were each divided in two, with all neurons recorded in the Rgb included in one data set and all neurons recorded in the RSA included in the other. Distance and classifier calculations were then computed independently for each regional data set in the same manner as described for the complete data set. Comparisons between regions and factors of interest were conducted using a two-way ANOVA.

Acknowledgements

I am particularly grateful for the contributions of Anna Serrichio, who completed her undergraduate thesis while assisting on this project. I also thank Alexandra Tse and Chuning Shi for assistance with animal training and electrode hyperdrive fabrication.

References

Alexander, A. S., & Nitz, D. A. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nat Neurosci, 18*(8), 1143-1151. doi:10.1038/nn.4058

Alexander, A. S., & Nitz, D. A. (2017). Spatially Periodic Activation Patterns of Retrosplenial Cortex Encode Route Sub-spaces and Distance Traveled. *Current Biology*. doi:10.1016/j.cub.2017.04.036

Auger, S. D., Mullally, S. L., & Maguire, E. A. (2012). Retrosplenial cortex codes for permanent landmarks. *PLoS One, 7*(8), e43620. doi:10.1371/journal.pone.0043620

Auger, S. D., Zeidman, P., & Maguire, E. A. (2015). A central role for the retrosplenial cortex in de novo environmental learning [JOUR]. *eLife, 4*, e09031. doi:10.7554/eLife.09031

Bar, M., & Aminoff, E. (2003). Cortical analysis of visual context. *Neuron, 38*(2), 347-358.

Bowers, D., Verfaellie, M., Valenstein, E., & Heilman, K. M. (1988). Impaired acquisition of temporal information in retrosplenial amnesia. *Brain Cogn, 8*(1), 47-66.

Cho, J., & Sharp, P. E. (2001). Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behav Neurosci, 115*(1), 3-25.

Cowansage, K. K., Shuman, T., Dillingham, B. C., Chang, A., Golshani, P., & Mayford, M. (2014). Direct reactivation of a coherent neocortical memory of context. *Neuron, 84*(2), 432-441. doi:10.1016/j.neuron.2014.09.022

Eichenbaum, H. (2014). Time cells in the hippocampus: a new dimension for mapping memories. *Nat Rev Neurosci, 15*(11), 732-744. doi:10.1038/nrn3827

Hyman, J. M., Ma, L., Balaguer-Ballester, E., Durstewitz, D., & Seamans, J. K. (2012). Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proc Natl Acad Sci U S A, 109*(13), 5086-5091. doi:10.1073/pnas.1114415109

Jezek, K., Henriksen, E. J., Treves, A., Moser, E. I., & Moser, M. B. (2011). Theta-paced flickering between place-cell maps in the hippocampus. *Nature, 478*(7368), 246-249. doi:10.1038/nature10439

Katche, C., Bekinschtein, P., Slipczuk, L., Goldin, A., Izquierdo, I. A., Cammarota, M., & Medina, J. H. (2010). Delayed wave of c-Fos expression in the dorsal hippocampus involved specifically in persistence of long-term memory storage. *Proc Natl Acad Sci U S A, 107*(1), 349-354. doi:10.1073/pnas.0912931107

Katche, C., Dorman, G., Gonzalez, C., Kramar, C. P., Slipczuk, L., Rossato, J. I., . . . Medina, J. H. (2013). On the role of retrosplenial cortex in long-lasting memory storage. *Hippocampus, 23*(4), 295-302. doi:10.1002/hipo.22092

Keene, C. S., & Bucci, D. J. (2008a). Involvement of the retrosplenial cortex in processing multiple conditioned stimuli. *Behav Neurosci, 122*(3), 651-658. doi:10.1037/0735-7044.122.3.651

Keene, C. S., & Bucci, D. J. (2008b). Neurotoxic lesions of retrosplenial cortex disrupt signaled and unsignaled contextual fear conditioning. *Behav Neurosci, 122*(5), 1070-1077. doi:10.1037/a0012895

Law, L. M., Bulkin, D. A., & Smith, D. M. (2016). Slow stabilization of concurrently acquired hippocampal context representations. *Hippocampus, 26*(12), 1560-1569. doi:10.1002/hipo.22656

Liu, X., Ramirez, S., Pang, P. T., Puryear, C. B., Govindarajan, A., Deisseroth, K., & Tonegawa, S. (2012). Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature, 484*(7394), 381-385. doi:10.1038/nature11028

Manns, J. R., Howard, M. W., & Eichenbaum, H. (2007). Gradual changes in hippocampal activity support remembering the order of events. *Neuron, 56*(3), 530-540. doi:10.1016/j.neuron.2007.08.017

Miller, A. M. P., Vedder, L. C., Law, L. M., & Smith, D. M. (2014). Cues, context, and long-term memory: the role of the retrosplenial cortex in spatial cognition [Review]. *Frontiers in Human Neuroscience, 8*. doi:10.3389/fnhum.2014.00586

Robinson, S., Keene, C. S., Iaccarino, H. F., Duan, D., & Bucci, D. J. (2011). Involvement of retrosplenial cortex in forming associations between multiple sensory stimuli. *Behav Neurosci, 125*(4), 578-587. doi:10.1037/a0024262

Smith, D. M., & Bulkin, D. A. (2014). The form and function of hippocampal context representations. *Neurosci Biobehav Rev, 40c*, 52-61. doi:10.1016/j.neubiorev.2014.01.005

Smith, S. M. (1979). Remembering in and out of context. *Journal of Experimental Psychology: Human Learning and Memory, 5*(5), 460-471. doi:10.1037/0278-7393.5.5.460

Spreng, R. N., Mar, R. A., & Kim, A. S. (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*(3), 489-510. doi:10.1162/jocn.2008.21029

Takahashi, N., Kawamura, M., Shiota, J., Kasahata, N., & Hirayama, K. (1997). Pure topographic disorientation due to right retrosplenial lesion. *Neurology, 49*(2), 464-469.

Todd, T. P., Meyer, H. C., & Bucci, D. J. (2015). Contribution of the retrosplenial cortex to temporal discrimination learning. *Hippocampus, 25*(2), 137-141. doi:10.1002/hipo.22385

van Groen, T., & Wyss, J. M. (1990). Connections of the retrosplenial granular a cortex in the rat. *J Comp Neurol, 300*(4), 593-606. doi:10.1002/cne.903000412

van Groen, T., & Wyss, J. M. (1992). Connections of the retrosplenial dysgranular cortex in the rat. *J Comp Neurol, 315*(2), 200-216. doi:10.1002/cne.903150207

van Groen, T., & Wyss, J. M. (2003). Connections of the retrosplenial granular b cortex in the rat. *J Comp Neurol, 463*(3), 249-263. doi:10.1002/cne.10757

Vedder, L. C., Miller, A. M., Harrison, M. B., & Smith, D. M. (2016). Retrosplenial Cortical Neurons Encode Navigational Cues, Trajectories and Reward Locations During Goal Directed Navigation. *Cereb Cortex*. doi:10.1093/cercor/bhw192

Wheeler, A. L., Teixeira, C. M., Wang, A. H., Xiong, X., Kovacevic, N., Lerch, J. P., . . . Frankland, P. W. (2013). Identification of a functional connectome for long-term fear memory in mice. *PLoS Comput Biol, 9*(1), e1002853. doi:10.1371/journal.pcbi.1002853

Figures

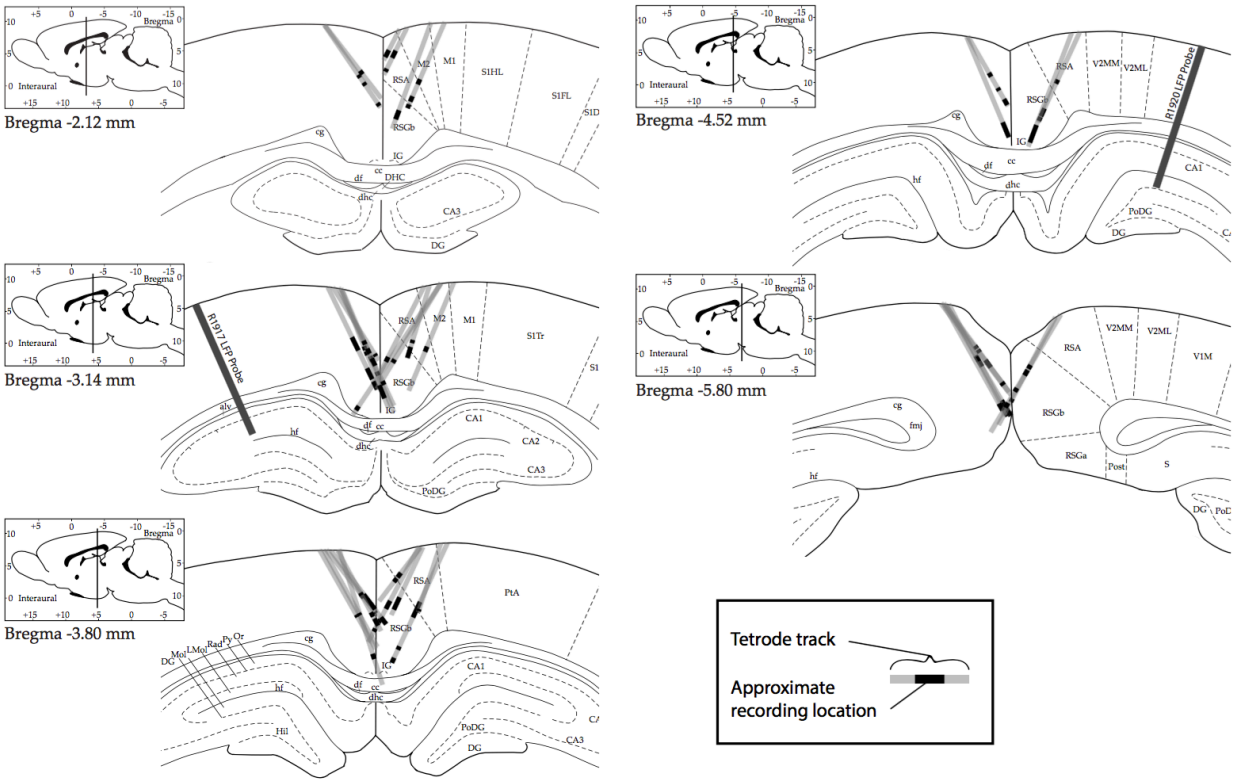

Figure 9 Estimated tetrode positions during recording (gray lines) and the approximate recording locations (overlaid black lines) are shown at 5 AP positions. Neurons were recorded primarily in Rgb, although a sizeable minority was recorded in RSA. Any recordings from tetrodes located outside of the RSC were omitted from the analysis. The locations of two local field potential probes that were not used in this study are also shown.

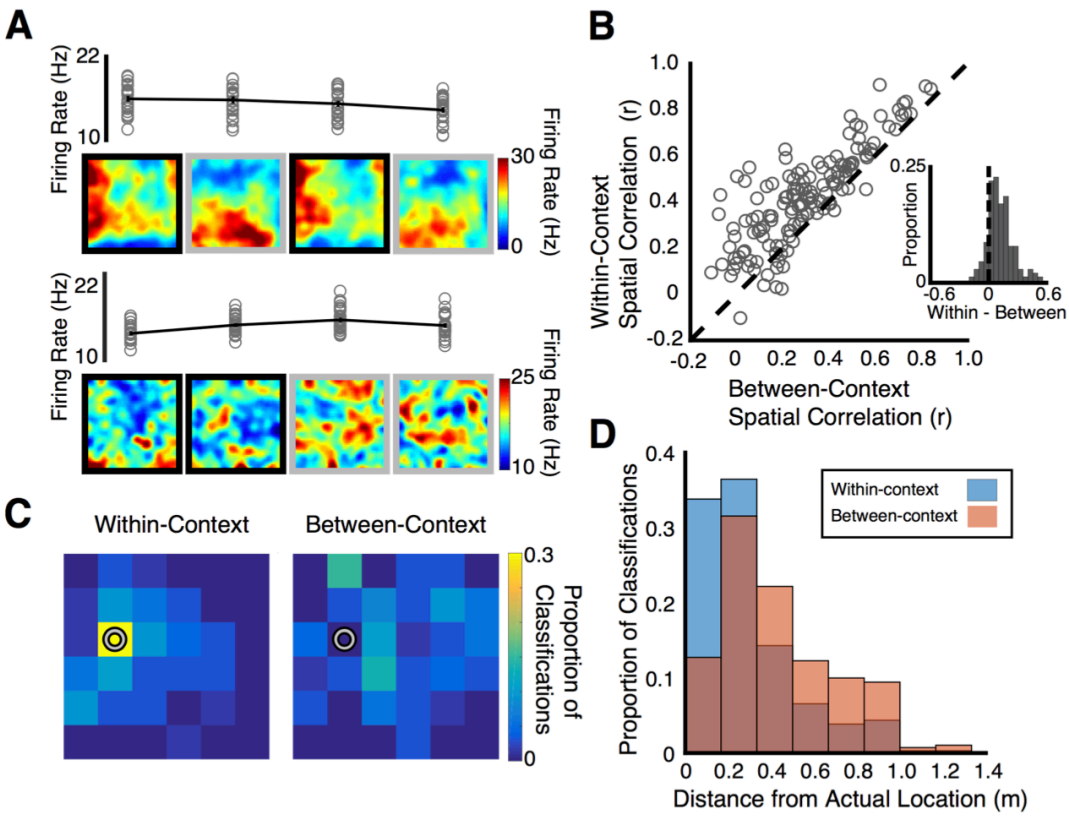

Figure 10 RSC ensembles show context-specific spatial coding. (**A**) Plots of firing rates and firing rate maps from two neurons during visits to each of the four contexts. These neurons exemplify clear (top) or typical (bottom) context-specific spatial coding. Open circles show firing rates during each 30 s time window, and black lines show mean firing rates with error bars showing +/- SEM. Squares around firing rate maps indicate the context during that visit. Note how the firing rate maps are more similar within context type than between context types. (**B**) Pixel-by-pixel correlations between context firing rate maps for all cells are plotted as open circles. Circles to the left of the unity line showed more similar spatial coding between two visits to the same context than between two different contexts. A clear majority of neurons showed greater correlations within contexts (inset). (**C**) Spatial classifications of all time windows from a single, representative spatial location are shown for a classifier trained on time windows recorded in either the same context (left) or the other context (right). The grey circle indicates the

actual location of the rats. Note that spatial classification is more accurate (greater proportion of the classifications are near to the actual location) with within-context classification. (**D**) The proximity of all classifications to the actual location of the rat is shown for within-context and between-context classification.

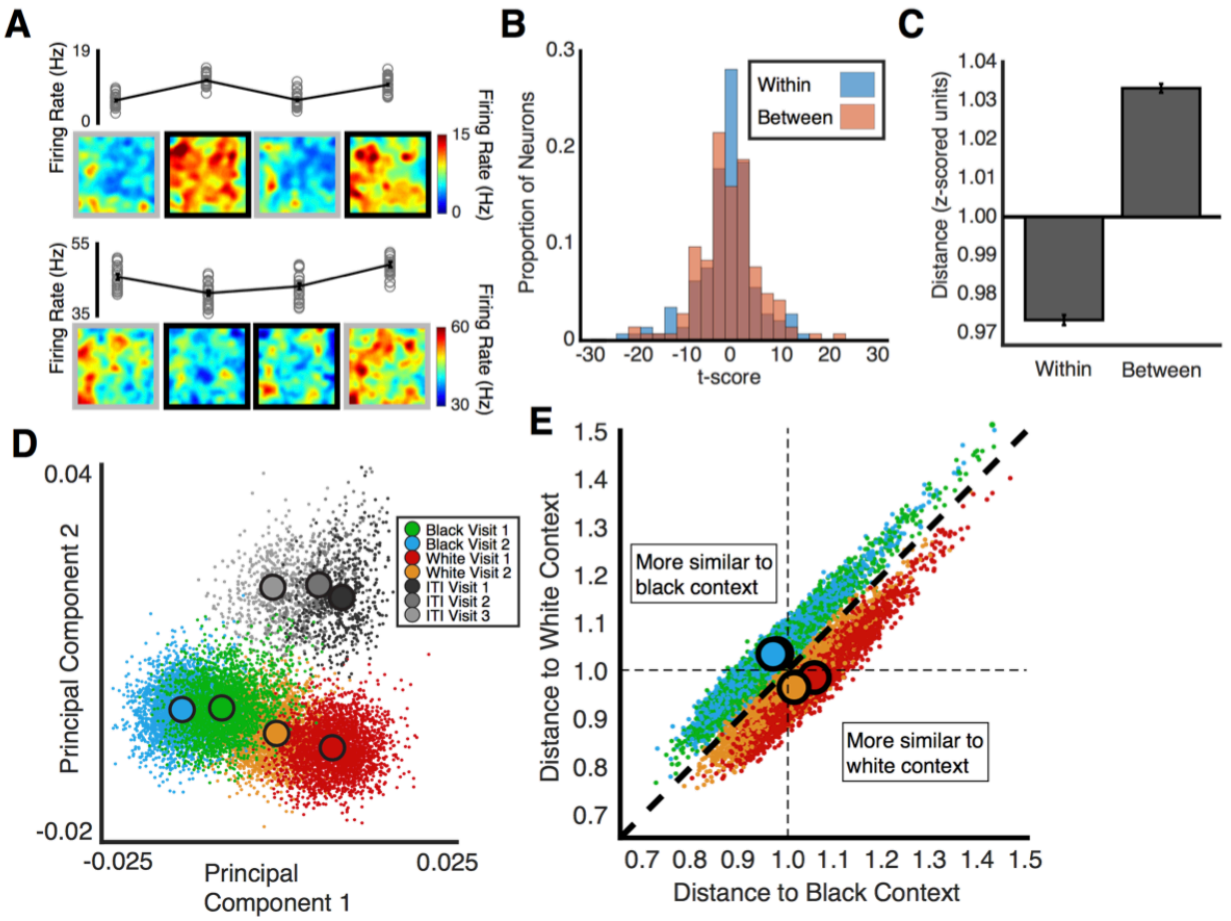

Figure 11 RSC ensembles uniquely encode different environments. (**A**) Plots of firing rates and firing rate maps from two neurons during visits to each of the four contexts. These neurons exemplify clear (top) or typical (bottom) context-specific coding. Open circles show firing rates during each 30 s time window, and black lines show mean firing rates with error bars showing +/- SEM. Squares around firing rate maps indicate the context during that visit. Note how firing rates are more similar within context type than between context type. (**B**) Histogram illustrating the distribution of context-specific firing in individual RSC neurons. For each neuron, we calculated the difference between the mean firing rate of all 30s time windows in the first and second visit to each context, and then divided this by the pooled standard deviation (i.e. the Z-score difference; within, red). We did the same for black and white contexts (between, red). (**C**)

Standardized Euclidean distances between RSC ensemble activity during two visits to the same context (within) and between visits to different contexts (between) is shown. Ensemble activity was far more similar within, compared to between, contexts. (**D**) The first two principal components are shown for activity during 250 ms time windows over the entire session. Black circles show context or ITI means. Note how the time windows cluster into their context or ITI of origin (colors). (**E**) Colored dots show ensemble activity from each time window plotted in terms of its distance from the mean of the same and other context representations. Points along the unity line are equidistant to both contexts, while points farther from the line indicate stronger ensemble preferences for one context over the other. Large dots illustrate the mean for each context visit. Note that ensemble activity reliably prefers the same context to the opposite.

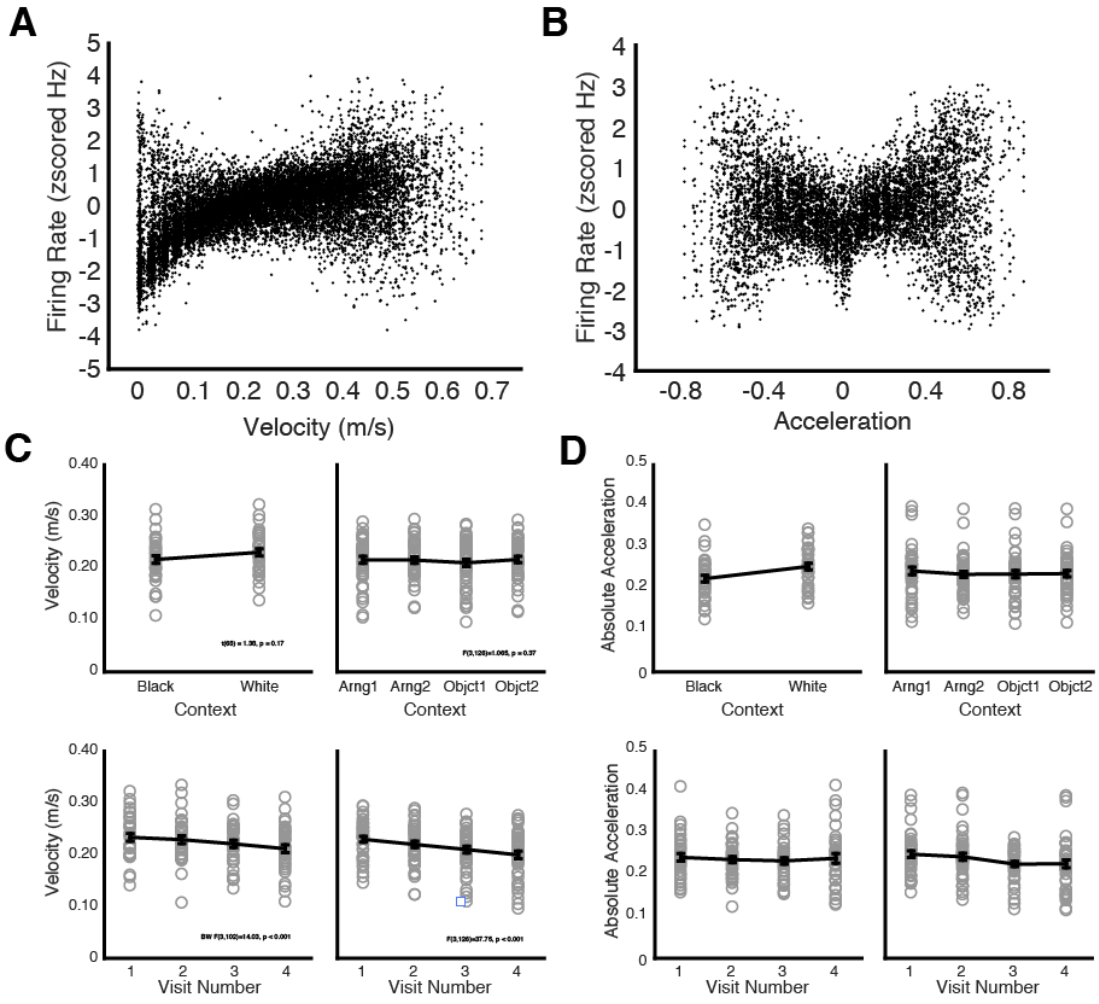

Figure 12 RSC neural firing is correlated with velocity and acceleration. (**A**) Relationship between RSC firing rate and running velocity is shown for all neurons. Average firing rate was calculated for 30 evenly spaced velocity windows (e.g., 0.1-0.15 m/s) spanning the full range of velocities exhibited by the rat during that session. Dots show the average normalized firing rate observed at that velocity. The relationship between velocity and firing rate appears more variable at the ends of the range due to fewer samples at those velocities (i.e. rats rarely ran faster than 0.5 m/s). (**B**) Relationship between RSC firing rate and acceleration is shown for all neurons, plotted as in A. (**C**) Relationship between velocity and context (top), and velocity and visit (bottom), is shown for every recording session. Open circles show the average velocity while the rat was in

that context or context visit. Solid lines show average velocity +/- SEM. Rats showed similar average velocities in black and white contexts ($t(68) = 1.38$, $p = 0.17$), and in each of the four object conditions ($F(3,126) = 1.07$, $p = 0.37$). However, velocity decreased over the four visits in both the black-and-white contexts ($F(3,102) = 14.03$, $p < 0.001$) and in the object conditions ($F(3,126) = 37.75$, $p < 0.001$). (**D**) Relationship between acceleration and context is show, plotted as in C. Rats showed greater average acceleration in the white context than in the black context ($t(68) = 2.63$, $p = 0.014$), but showed similar acceleration in the four object conditions ($F(3, 126) = 0.51$, $p = 0.68$). Rats showed similar acceleration across the four visits to the black and white contexts ($F(3,102) = 0.42$, $p = 0.74$). However, acceleration decreased over the four visits to the object conditions ($F(3,126) = 5.982$, $p < 0.01$).

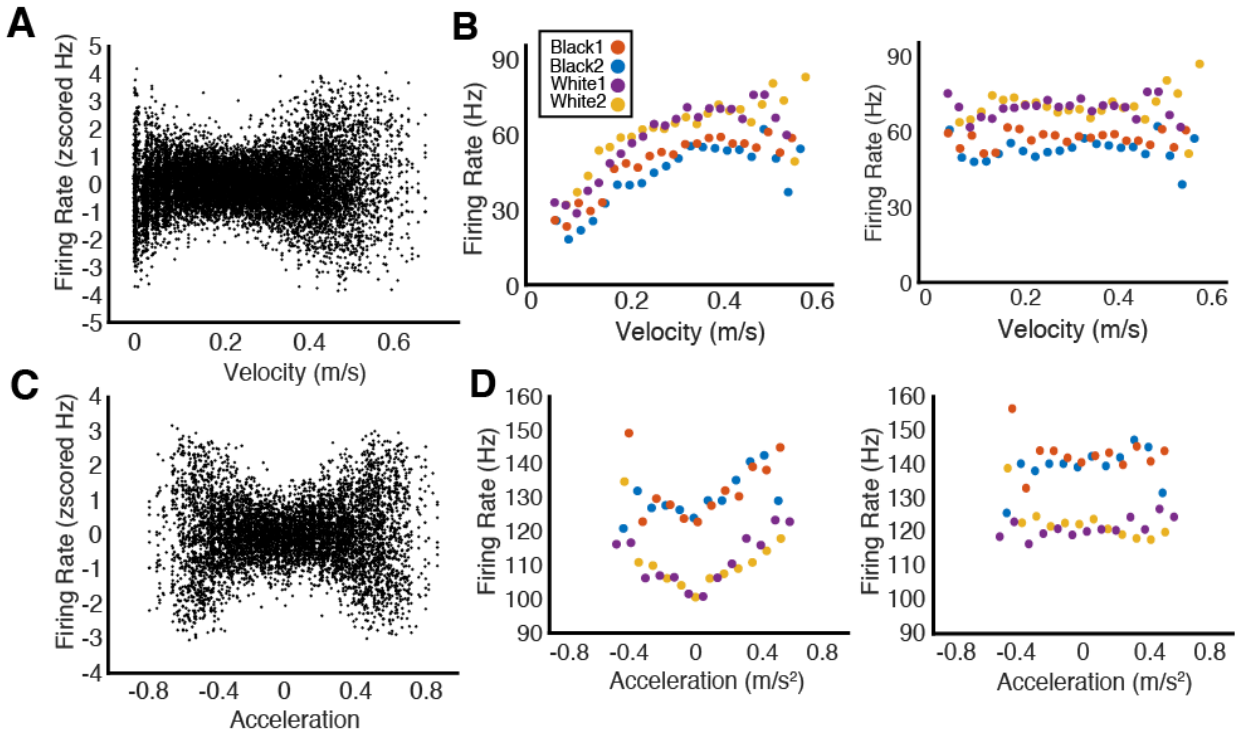

Figure 13 Controlling for relationships between firing rate and movement variables. (**A**) Relationship between RSC firing rate and running velocity is shown for all neurons after statistically correcting for the relationship between the two variables (see methods). Dots indicate the average normalized firing rate observed at that velocity. (**B**) Example neuron showing the effects of the statistical correction for velocity. Originally firing rate increased with velocity (left). After the correction, firing rate remained constant over the range of velocities (right). Note that this cell fired more during visits to the white context, and that this was true both before and after the correction. (**C**) Corrected relationship between RSC firing rate and acceleration is shown for all neurons, plotted as in A. (**D**) Example neuron showing the effects of the statistical correction for acceleration. Originally firing rate increased with both acceleration and deceleration (left). After the correction, firing rate remained constant over the range of

accelerations (right). Note that this cell fired more during visits to the black context, and that this was true both before and after the correction.

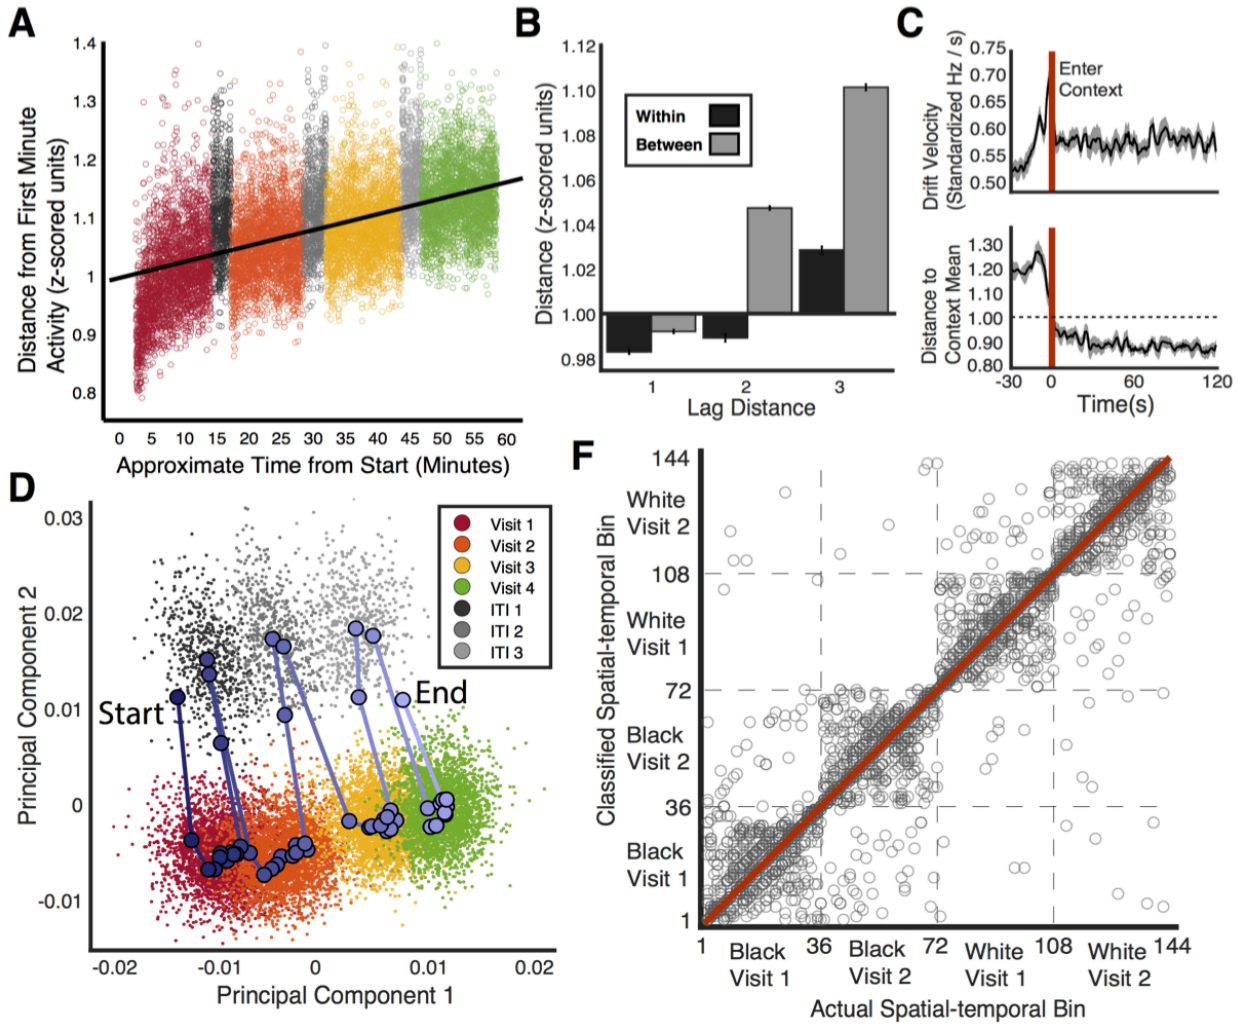

Figure 14 Neural populations in the RSC encode visit order. (**A**) Colored circles show ensemble activity from each time window plotted in terms of its distance from the mean of ensemble activity during the first minute of the session. Note how the ensemble moves steadily away from its starting state over the course of the session. (**B**) Standardized Euclidean distance between RSC ensemble activity occurring during two visits to the same (black) or two different (grey) context visits is plotted in terms of how far the two visits were apart in their presentation order during the session. For example, back-to-back visits were 1 lag distance apart, while visits at the beginning and the end of the session were 3 apart. (**C**) Average ensemble activity is plotted over

the window from 30 s before to 120 s after the rats' entrance into each context. The distance between ensemble activity during adjacent time windows divided by time (i.e. drift velocity) peaked immediately before rats entered the context and then leveled off (top). The distance from ensemble activity to the mean of ensemble activity in that context decreased rapidly as rats entered the context and then leveled off (bottom). (**D**) The first two principal components are shown for activity during 250 ms time windows over the entire session. Blue lines trace the state of the RSC ensemble between means of activity calculated at each minute (blue circles). Lighter blue colors indicate means occurring later in the session. Note how the smooth drift in the state of the RSC ensemble (from left to right) is accompanied by sharp shifts (up and down) between context and ITI visits.

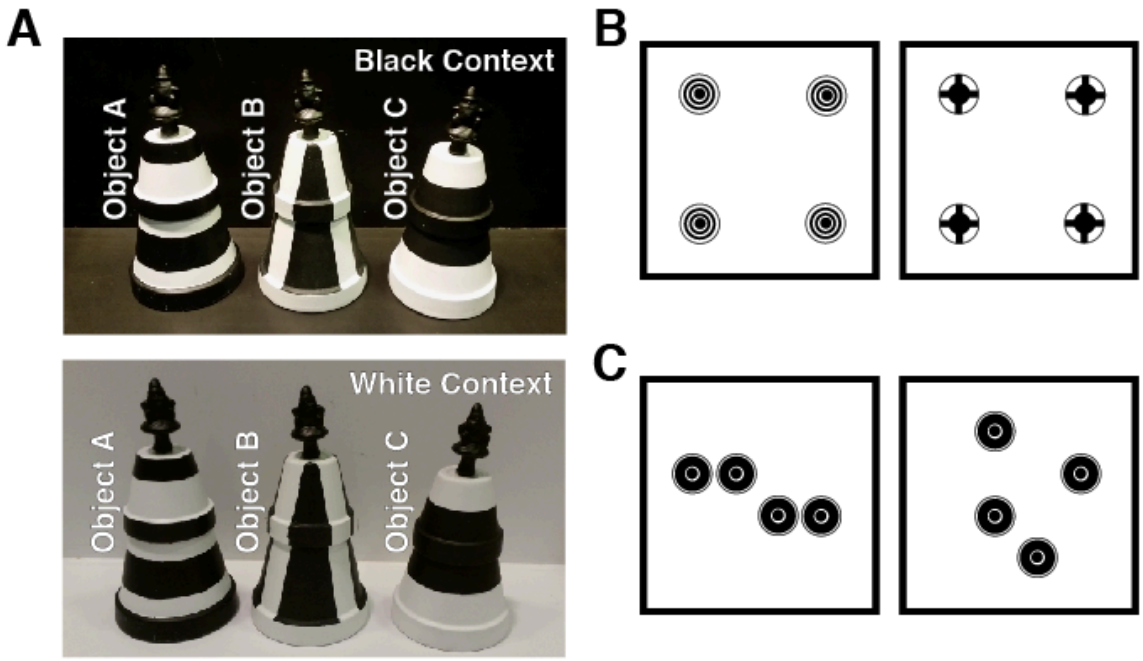

Figure 15 Object appearance and arrangement stimuli and manipulations. (**A**) The three distinct objects (A, B, C) are shown in each of the two contexts (Black and White). (**B**) The object appearance manipulation employed objects A (left) and B (right) presented in the same arrangement. (**C**) For the object arrangement manipulation, object C was presented in one of two distinct arrangements.

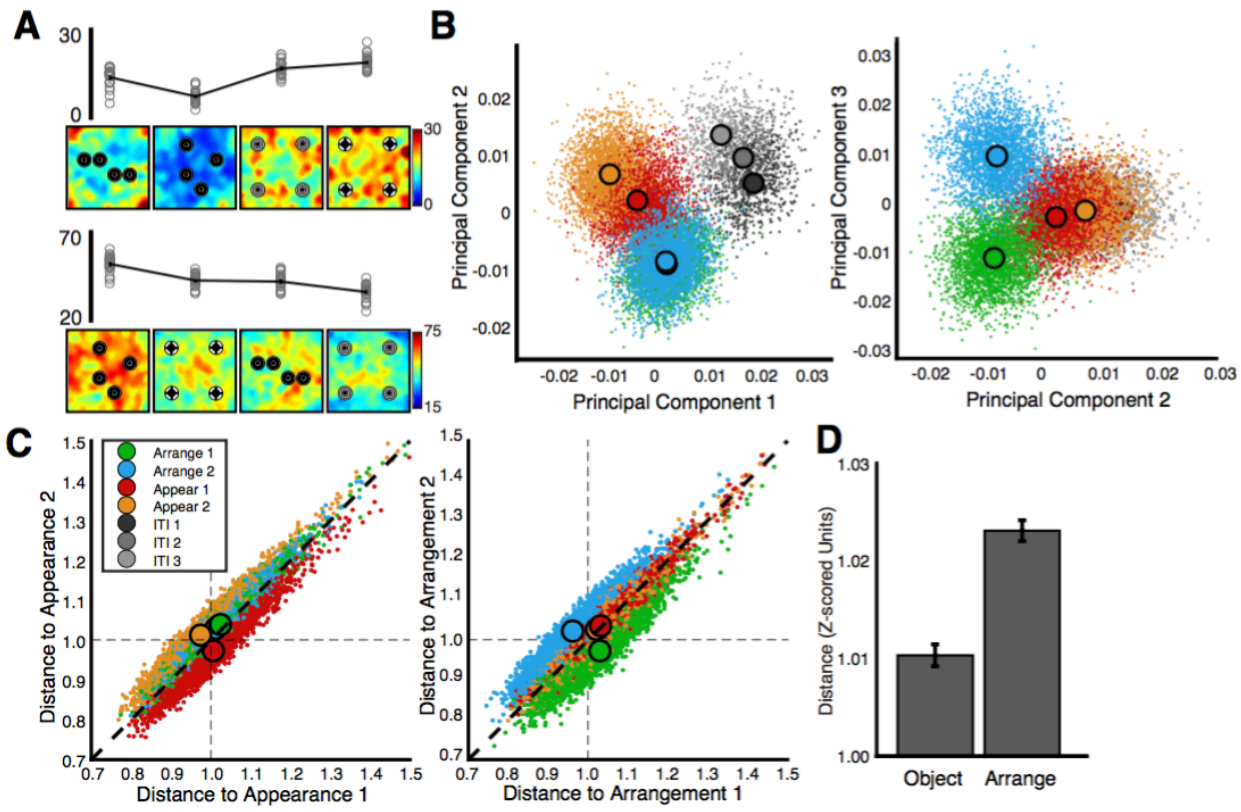

Figure 16 RSC ensembles preferentially encode spatial arrangements. (**A**) Plots of firing rates and firing rate maps from two representative neurons during visits to each of the four object conditions. Open circles show firing rates during each 30 s time window, and black lines show mean firing rates with error bars showing +/- SEM. (**B**) The first three principal components are shown in two plots for activity during 250 ms time windows over the entire session. (**C**) Colored dots show ensemble activity from each time window plotted in terms of its distance from the mean of the two object appearance conditions (left) and the two object arrangement conditions (right). Points along the unity lines are equidistant to both conditions, while points farther from the line indicate stronger ensemble preferences for one condition over the other. Large dots illustrate the mean for each condition. Note that time windows from the arrangement conditions are further off the unity line (i.e. show a greater preference for their condition of origin) than are

time windows from the appearance conditions. (**D**) Standardized Euclidean distances between RSC ensemble activity during the two visits to the object appearance conditions (appear) and between the two visits to the object arrangement conditions (arrange) is shown. Ensemble activity was more distinct between visits to the arrangement conditions.

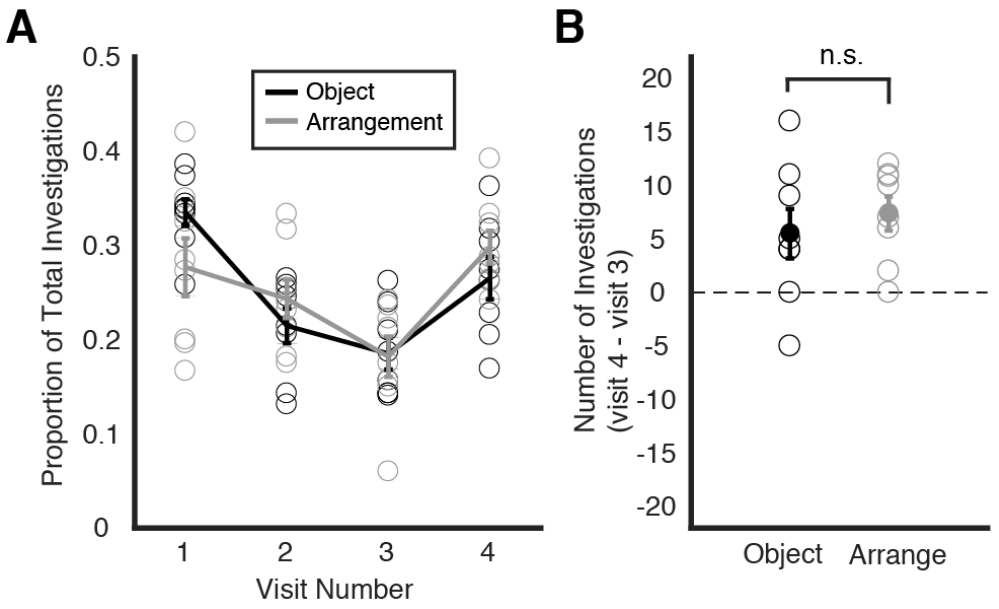

Figure 17 Object appearance and arrangement manipulations are similarly salient. Rats made four visits to the same context (Black Context or White Context) where they explored four copies of the same object presented in a specific spatial arrangement. The first three visits repeated the same four objects presented in the same spatial arrangement. Then, to test whether rats could detect changes to the presentation, either the appearance of the objects (different objects were used) or the spatial arrangement of the objects (objects were moved to different locations) was changed for the fourth visit. The rat's interest in the fourth (novel) presentation was measured in terms of number of object investigations. (**A**) The proportion of object investigations over the four visits is shown for all rats. Open circles show data from one rat from one visit. Solid lines show means for all rats +/- SEM. Note that the number of investigations declined over the three repeat visits, and then increased after the object or arrangement was changed. (**B**) Rats in both the object and arrangement conditions showed a significant increase in the number of investigations from the third to fourth visit, and there was no difference in the size of the increase between the two conditions. A two-way repeated measures ANOVA comparing the number of investigations on days 3 and 4 revealed a main effect of day ($F(1,14) = 21.20$, $p <$

0.001), but no affect of condition ($F(1,14) = 0.94$, $p = 0.350$, and no day by condition interaction ($F(1,14) = 0.45$, $p = 0.513$).

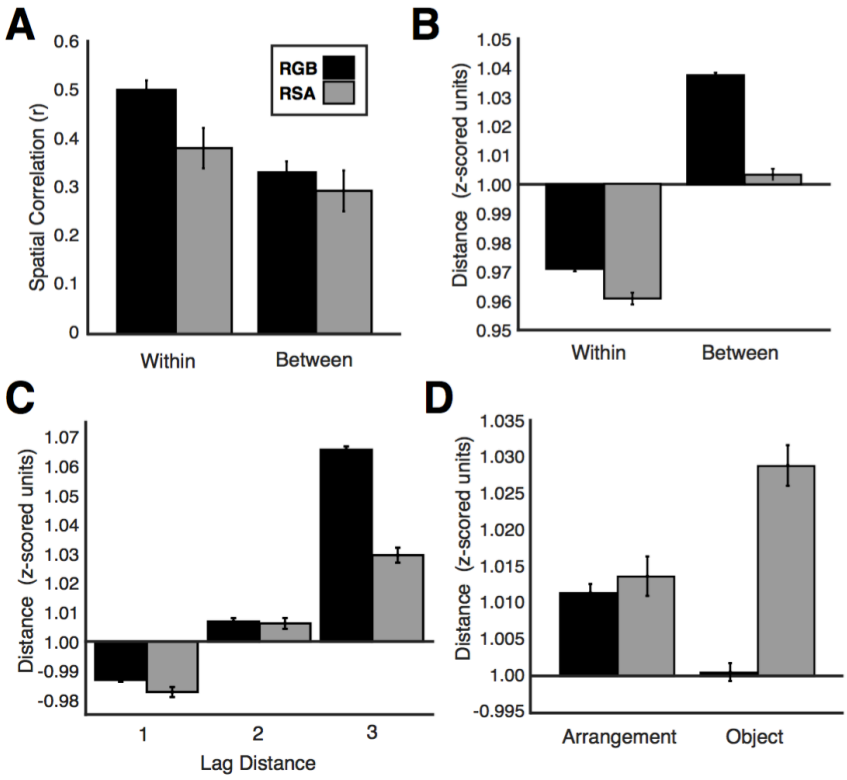

Figure 18 Spatial and temporal coding in RSC subregions. (**A**) Neurons in both the Rgb and RSA subregions showed greater spatial correlations between two visits to the same context (within) than between two visits to different contexts (between). Bars show means +/- SEM. (**B**) Standardized Euclidian distances between RSC ensemble activity within-contexts was far less than between-contexts in both Rgb and RSA. However, Rgb showed far greater distances between contexts than did RSA. (**C**) Distances between RSC ensemble within-context activity increased with the lag distance between visits in both Rgb and RSA. (**D**) In the Rgb, distances between RSC ensemble activity occurring during visits to the two arrangement conditions were greater than between visits to the two object conditions. However, the opposite was true in the RSA.

CHAPTER THREE

Retrosplenial Coding of Long-term Spatial Memories and Goal Locations

Abstract: The RSC plays a key role in spatial and contextual memory. However, surprisingly little is known about spatial representations in the RSC or how such representations might emerge as a result of experience. We trained rats on a continuous spatial alternation task and recorded RSC neuronal activity throughout learning. We found that the RSC developed a rich population-level representation of the rat's spatial location, the goal locations, and trajectories to the goal. In addition to the rat's current location, RSC ensembles also represented the goal locations as the rats approached the choice point. After learning, when the rats reached peak performance, these neural simulations became selective for the correct future reward location. Many of these representations emerged slowly as the rats learned the task, suggesting that the RSC encodes LTM that can be used to support memory-guided spatial navigation.

Introduction

LTM is thought to rely on the neocortex and is critical for a variety of cognitive processes, including attention, decision making, and new learning (Buckner & Carroll, 2007; Chun, 2000; Tse et al., 2007). In spatial navigation, for example, subjects use LTM for the spatial layout of the environment, including routes, goal locations, and task rules, to achieve their navigation goals (Tolman, 1948). Much of what we know about spatial memory and navigation comes from work on the hippocampus (Buzsaki & Moser, 2013). However, many spatial memories do not require the hippocampus after they have become well-learned (Day, Langston, & Morris, 2003; Maguire, Nannery, & Spiers, 2006; Winocur, Moscovitch, Fogel, Rosenbaum, & Sekeres, 2005), and several studies have suggested that spatial and contextual LTM is likely stored in a network of cortical regions, including prefrontal, anterior cingulate, and RSC areas (Maviel, Durkin, Menzaghi, & Bontempi, 2004; Tse et al., 2011; Wheeler et al., 2013).

The RSC, in particular, appears to contribute to the storage of spatial LTM (Miller, Vedder, Law, & Smith, 2014). The RSC is anatomically interconnected with brain regions known to play a role in navigation, including the hippocampus and the ATN (van Groen & Wyss, 1992, 2003). Recent findings suggest that contextual LTM depends on the RSC (Corcoran et al., 2011; Cowansage et al., 2014; Katche et al., 2013), and that RSC damage impairs navigation in humans and rodents (Czajkowski et al., 2014; Ino et al., 2007; Pothuizen, Davies, Aggleton, & Vann, 2010; Takahashi, Kawamura, Shiota, Kasahata, & Hirayama, 1997). An extensive fMRI literature similarly suggests an RSC role in spatial memory (Auger, Zeidman, & Maguire, 2015; Morgan, Macevoy, Aguirre, & Epstein, 2011; Wolbers & Buchel, 2005), including in the representation of spatial heading (Marchette, Vass, Ryan, & Epstein, 2014; Shine, Valdes-Herrera, Hegarty, & Wolbers, 2016) and landmarks (Auger & Maguire, 2013;

Auger, Mullally, & Maguire, 2012). Furthermore, individual RSC neurons exhibit spatially localized firing (Cho & Sharp, 2001; Smith, Barredo, & Mizumori, 2012) and directional firing (Chen, Lin, Green, Barnes, & McNaughton, 1994; Cho & Sharp, 2001; Jacob et al., 2017), and they respond strongly to navigational cues (Vedder, Miller, Harrison, & Smith, 2016). Recent work has also shown that RSC neuronal activity is modulated by allocentric, egocentric, and route-centered spatial reference frames (Alexander & Nitz, 2015, 2017), consistent with the idea that the RSC is well positioned to integrate information from each of these domains (Vann, Aggleton, & Maguire, 2009).

However, much remains unknown about the RSC role in navigation, such as whether RSC neurons generate a map-like representation that indicates the subject's position in the environment and, if so, how such information is encoded, whether it develops as a result of repeated experiences, or how these representations might be used to solve spatial memory tasks. To answer these questions, we trained rats on a continuous spatial alternation task and recorded RSC neuronal activity throughout learning.

Results

RSC lesions impair spatial navigation at asymptote.

Although the RSC is important for many spatial tasks in rats, including some alternation tasks (Nelson, Powell, Holmes, Vann, & Aggleton, 2015; Pothuizen et al., 2010), it is not known whether the RSC is needed for continuous spatial alternation, which is widely used in neurophysiology research. We evaluated this by training control rats and rats with neurotoxic (NMDA) lesions of the RSC on a modified t-maze designed for continuous spatial alternation (**Figure 19A**). The rats were rewarded with 0.2 ml of chocolate milk for alternating between the left and right reward arms of the maze. Rats were given daily training sessions until they reached a criterion of 90% correct, followed by four additional asymptotic performance sessions (**Figure 20A**). Because rats required varying numbers of sessions to reach the criterion, we restricted our analyses to a set of training stages common to all subjects, including the first session, the middle training session (half way between the first session and asymptote), the session in which the rats reached the criterion, and the four post-criterial overtraining sessions. A two-way repeated measures ANOVA comparing control and lesion group performance during these learning stages revealed a significant main effect of training stage ($F_{(6,90)} = 62.66$, $p < 0.001$) and a significant training stage by lesion group interaction ($F_{(6,90)} = 2.832$, $p < 0.05$, **Figure 19B**). Post-hoc comparisons showed that there was no effect of RSC lesions during acquisition (i.e. the first, middle, and criterial sessions; $t_{(15)} = 0.33$, $p = 0.75$; **Figure 20B top**). There was also no lesion impairment in terms of the number of sessions needed to reach the criterion ($t_{(15)} = 0.19$, $p = 0.85$; **Figure 20C top**). However, there was a modest but highly reliable impairment during asymptotic performance ($t_{(15)} = 3.47$, $p < 0.005$; **Figure 19B inset left**), consistent with findings

suggesting that the RSC is preferentially involved in the late stages of learning (Bussey, Muir, Everitt, & Robbins, 1996; Gabriel, 1993; Wolbers & Buchel, 2005) and LTM (Katche et al., 2013). The size of the lesions varied considerably (39.5 - 72.8% of the RSC, **Figure 20D,E**), and the impairment in asymptotic performance was tightly correlated with lesion size ($r = -0.90$, $p < 0.01$; **Figure 19B inset right**). Performance during learning ($r = -0.36$, $p = 0.43$) and the number of sessions needed to reach the criterion ($r = 0.38$, $p = 0.40$; **Figure 20B,C bottom**) were not correlated with lesion size.

RSC neural populations develop a spatial representation of the maze with learning

In order to examine the neurophysiological basis of the RSC role in spatial cognition, we recorded from 637 RSC neurons in 12 rats as they learned and performed the continuous spatial alternation task (336 neurons during acquisition and 301 during asymptotic performance). Our recordings targeted Rgb subregion of the RSC bilaterally, although small numbers of neurons from Rga and RSA were also included. There were no conspicuous differences in the firing properties of neurons recorded in different subregions, different hemispheres, or at different AP coordinates (see methods and **Figure 21** for anatomical details). Previous studies have shown that RSC neurons exhibit spatially localized firing (Cho & Sharp, 2001). However, more recent studies suggest that spatial firing in the RSC is less specific than in the hippocampus (Alexander & Nitz, 2015; Smith et al., 2012). Consistent with this, we found that although 18.64% of the RSC neurons of the present study could be classified as place cells with previously used criteria (Smith et al., 2012) (**Figure 19C,D**, also see **Figure 22A**), they exhibited larger fields, higher background firing rates, and lower spatial information content than those reported in the hippocampus (Ito, Zhang, Witter, Moser, & Moser, 2015; Smith et al., 2012). The distribution of

the spatial information content scores suggests that these RSC place cells were simply extreme cases from a population that ranged from very little to moderate levels of spatial selectivity, rather than from a distinct subset of RSC neurons that exhibited true place fields. Instead, RSC neurons tended to fire at high rates across large areas of the maze (**Figure 22B**).

Despite the relatively modest spatial coding of individual RSC neurons, the RSC developed a detailed and reliable spatial representation at the population level as the rats learned the alternation task. To quantify this, we took a Bayesian decoding approach to determine whether we could predict the rat's current spatial location solely on the basis of the activity of the recorded neural ensemble (Wilson & McNaughton, 1993; Zhang, Ginzburg, McNaughton, & Sejnowski, 1998) (**Figure 19E**). We found that the rat's position could be accurately predicted at a rate far greater than chance even on the first day of learning ($p < 0.001$, compared to a control distribution generated by shuffling firing rates across time bins, see methods; **Figure 19F**), and that decoding accuracy improved significantly as the rats learned ($F(3, 30) = 3.02$, $p < 0.05$). At asymptote, the rat's position could be accurately predicted within 4.5 cm about 40% of the time.

To confirm that this improvement in spatial representation could not have been a product of learning-related changes in behavior (e.g., more consistent trajectories around the maze later in learning), we employed an additional correlational reconstruction technique that limited the analysis to trials in which the rat followed a direct and stereotyped path through each maze area. To do this, we combined neurons from all rats into a single population for each training stage, and calculated mean firing rate vectors for each spatial bin along the paths (**Figure 19G**), separately for the first and second halves of each session. Mean ensemble firing rate vectors from the first half of the session were then correlated with those from the second half. If spatially localized firing was reliable across session halves, the highest correlation should occur between

visits to the same spatial bin (during the first and second-half of the session), and therefore accurate spatial coding is indicated by peak correlations occurring along the diagonal of the correlation matrix (**Figure 19H**), whereas deviations from the diagonal indicate inconsistent or erroneous representations (i.e. spatial coding errors). Consistent with the Bayesian analysis above, spatial coding was far more reliable and accurate than expected by chance at all stages of learning (all $p < 0.001$, compared to a distribution generated by shuffling first-half/second-half neuron pairings) and the representation improved with learning, as indicated by a 70% reduction in spatial coding errors from early learning to overtraining ($p < 0.005$, compared to a distribution generated by shuffling neurons between stages; **Figure 19I**). This analysis included only a single stem traversal (from the right to left reward location) for simplicity because trial-type specific firing on left and right trials can affect the correlations. However, similar results were obtained when the stem was included twice (as two separate trajectories for left and right trials) or when the stem was excluded. The observation that two independent methods for assessing spatial representations, Bayesian decoding and correlational reconstruction, produced the same pattern of results provides the first compelling evidence that detailed spatial representations in the RSC develop with learning.

RSC neural populations represent each reward location

Goal locations are salient features of the environment, and we have previously shown that RSC neurons selectively respond to the co-occurrence of the reward and the specific location where it was obtained (Smith et al., 2012; Vedder et al., 2016). Reward locations may be particularly important for spatial alternation because future choice depends on memory for the previous reward. Using ANOVA as a classification tool, as in our previous work (Vedder et al.,

2016), we found here that 16.61% of individual RSC neurons exhibited a significant change in firing when the rat received the reward, specifically at one of the two (right or left) goal locations (**Figure 23A,B**). Significantly more location-specific reward responses were observed at every training stage than was expected from chance (all $p < 0.001$, compared to a distribution generated by shuffling left and right trial types), and this proportion did not change significantly over the course of training ($p = 0.33$, compared to a distribution generated by shuffling neurons between stages; **Figure 23B**). Some of these responses were tightly time-locked to the receipt of the chocolate milk reward (e.g., **Figure 23A top**), as determined by an automatic lick detection circuit. However, many other neurons were not time-locked, often including anticipatory changes in firing rate as the rat approached the reward, but nevertheless appeared to differentiate the two reward locations (e.g., **Figure 23A middle and bottom**).

This suggested that, as with the spatial coding described above, the encoding of goal locations may depend on a population code rather than on a subset of neurons with firing that was strictly driven by the reward. To examine this, we combined neurons from different rats into a single ensemble (without any *a priori* classification of reward-related firing) for each training stage and calculated the similarity between ensemble responses occurring at each of the two reward locations (see Methods). If ensemble firing patterns were highly specific to each reward location, then activity during visits to one reward location should be similar to other visits to that location, and dissimilar to visits to the opposite location. To calculate this, we used the standardized Euclidean distance between ensemble firing rate vectors as a measure of similarity (with small distances corresponding to greater similarity), and computed the distance from each individual reward visit to the mean firing rate vectors from each reward location (**Figure 23C,D**). Remarkably, the ensemble firing patterns distinguished the left and right reward

locations with perfect accuracy at every stage of learning (ensembles were more similar to the average activity of the same reward location than of the opposite location, see methods, **Figure 23D inset**). Similar classification results were obtained using linear discriminant analysis on a dimensionally reduced dataset (**Figure 24**). Furthermore, ensemble activity at the rewards became increasingly location specific with learning. Activity occurring at the left reward became more similar to other left reward visits, and less similar to right reward visits (and likewise with right rewards, $F(3,100) = 18.69$, $p < 0.001$, **Figure 23C**).

RSC neural populations develop trial-type specific responses on the stem of the maze

RSC neurons exhibited two distinct firing patterns on the stem of the maze depending on whether the rat was about to turn left or right, similar to previous reports on hippocampal neurons (Frank, Brown, & Wilson, 2000; Wood, Dudchenko, Robitsek, & Eichenbaum, 2000). Using the classification procedures of Wood and Eichenbaum (Wood et al., 2000), we found that firing on the stem became more specific to trial type (left or right) as the rats learned the task (**Figure 25A,B**). The proportion of neurons that met the criteria for trial-type specific firing on the stem during left and right trials increased significantly with training ($p < 0.05$, compared to a distribution generated by shuffling neurons between training stages). On the first and middle days of training, the proportion of trial-type specific neurons (10.13% and 11.49%, respectively) did not differ from chance (first, $p = 0.14$; middle, $p = 0.07$; both compared to distributions generated by shuffling left and right trial types). However, the proportion doubled by criterion day, far exceeding chance (20.77% on criterion day, $p < 0.001$; 21.02% during overtraining, $p < 0.001$). It is unlikely that trial-type specific firing resulted from small differences in movement

on the stem, as neither lateral position ($r = 0.08$, $p = 0.18$) nor running speed ($r = 0.08$, $p = 0.18$) was significantly correlated with trial-type specific firing (**Figure 26A-C**).

Although clearly visible in single neurons, trial-type specific firing patterns appeared to be present to varying degrees throughout the RSC population rather than in a particular subset of neurons (**Figure 26D**), suggesting that the trial-type specificity of stem firing might also be encoded at the population level. To examine this possibility, we combined neurons across subjects (as with the above analysis of ensemble reward responses) and used the resulting ensembles to predict upcoming left or right choices, and to determine whether greater trial-type specificity in the ensemble was associated with better alternation performance. Consistent with the single-unit analyses above, trial-type specificity increased with learning ($F(3,100) = 20.02$, $p < 0.001$, **Figure 25C,D**), indicating that ensemble activity on left trials became more similar to other left trials, and less similar to right trials (and likewise with right trials). To then determine how well we could predict choice behavior, we classified trials as either left or right depending on whether ensemble activity was more similar to the mean activity of left or of right trials, and found that classification accuracy increased with learning ($p < 0.05$, compared to a control distribution generated by shuffling neurons between training stages, **Figure 25E**), from chance classification on the first day (46.15%) to perfect classification during overtraining (100%). Another classification technique, linear discriminant analysis, provided qualitatively similar results (**Figure 26E**).

The finding that trial-type specificity increased with learning suggested that these representations might be related to alternation accuracy. To specifically assess this possibility, we binned the overtraining sessions into four categories based on performance, and found that sessions with better performance had higher trial-type specificity ($F(3, 132) = 24.38$, $p < 0.001$;

Figure 25F,G). Indeed, trial-type specificity was three-fold higher in the sessions with the best performance (96.1-100% accuracy) than in sessions with poor performance (84%-88% accuracy; $t(66) = 7.68$, $p < 0.001$). Similarly, our ability to classify trials as either left or right was better on superior performance days than on poor performance days ($p < 0.05$, compared to a distribution generated by shuffling neurons between performance groupings; **Figure 25H**). Thus, population representations of the left and right trajectories developed with training and were associated with better performance.

RSC populations represent future goal locations

In human subjects, the RSC is active during navigational route planning (Brown et al., 2016), and in rodents navigational decision-making is associated with forward-sweeping simulations of possible trajectories to the goal in the hippocampus (Johnson & Redish, 2007). Here, we used Bayesian decoding to examine the possibility that RSC ensembles encode information about future goal locations as the rats approached the choice point (**Figure 27A-C**). We computed probability distributions reflecting the predicted location of the rat given ensemble firing activity taken every 50 ms as the rats traversed the stem (see methods for details). Interestingly, we found that a sizeable portion (25.86%) of the trial-averaged probability distribution was located in the reward areas far ahead of the rat's actual location on the stem (reward areas are shown in **Figure 27B**), consistent with the idea that the RSC ensembles were representing the two goal locations. Furthermore, this decoding to the reward areas increased with learning ($F(3, 30) = 2.95$, $p = 0.05$, **Figure 27D**), from marginally greater than chance during the early and middle stages of learning (early, $t_{(5)}=2.22$, $p = 0.08$; middle, $t_{(4)} = 2.53$, $p =$

0.06), to far greater later in learning (late, $t_{(7)} = 4.76$, $p < 0.005$; overtraining, $t_{(14)} = 9.21$, $p < 0.001$).

While RSC ensembles initially represented both the correct and incorrect reward areas equally, they began to preferentially represent the correct reward area once the rats reached asymptote. To quantify this, we compared the decoded probability that the rat was located in the correct reward area, p(correct), with the opposite reward area, p(incorrect, **Figure 27B**), and found that RSC ensembles began to preferentially represent the correct reward area with learning ($F(3, 30) = 7.68$, $p < 0.001$, **Figure 27E**). Early in learning, RSC populations showed a numerical preference for the incorrect reward area. However, after correcting for multiple comparisons, we only found a significant preference between the two reward areas during overtraining sessions ($t_{(14)} = 3.36$, $p < 0.005$). This preference for the future reward area during overtraining was not correlated with lateral position on the stem ($r = 0.11$, $p = 0.69$; **Figure 27F top**) or with running speed ($r = 0.22$, $p = 0.43$; **Figure 27F bottom**). Similar results were obtained for all analyses using the proportion of absolute classified positions (as in **Figure 19E,F**) rather than the proportion of the decoded probability distribution.

Discussion

Neocortical regions such as the RSC play an important role in LTM storage, and recent work in humans and rodents has converged on the importance of the RSC for spatial memory. Our findings are the first direct evidence that a rich and detailed spatial representation develops in the RSC over the course of learning, including information about the rat's current spatial location, the current trajectory, reward locations, and even simulations of upcoming goal locations. In contrast to the rapid spatial coding seen in the hippocampus, RSC representations developed slowly as the rats learned the task and some representations, such as simulations of future goal locations, did not emerge until after learning. These results suggest that the RSC is a key component of the neocortical system for storing spatial LTM, consistent with recent molecular and optogenetic work (Cowansage et al., 2014; Czajkowski et al., 2014; Katche et al., 2013).

Despite relatively poor spatial specificity of individual neurons (Alexander & Nitz, 2015; Smith et al., 2012) (**Figure 19**, **Figure 22**) the RSC generates a surprisingly accurate population level representation of the rat's current spatial location. In well-trained rats, the current position could be decoded far above chance levels, even from ensembles with as few as eight neurons, and correlation analyses confirm that ensemble firing patterns at any given location are highly consistent across the session. Unlike in the hippocampus, where each position is encoded by a relatively small percentage of pyramidal neurons with highly specific place fields, most RSC neurons were active over the entire maze, with firing rates that were higher in some regions and lower in others. This pattern of results, with precise population level representations arising from relatively non-specific individual neurons, was seen throughout our data set.

RSC neurons did not exclusively encode the rat's current spatial position, but they were also sensitive to the rat's current trajectory, as indicated by firing patterns that were specific to the left and right trials as the rats traversed the stem. This is consistent with previous reports of directional sensitivity during a track shuttling task (Alexander & Nitz, 2015) and our previous finding of differential firing during left and right trials in a light-cued T-maze task (Vedder et al., 2016). Our results indicate that trial-type specific firing is also prevalent in a memory-guided navigation task. Although trial-type specificity could be seen in individual neurons, the firing patterns were distributed across the population and not confined to a specific subset of responsive neurons (**Figure 26D**). Our finding that these population-level representations were related to task performance (greater trial-type specificity was associated with greater choice accuracy), suggests that the memory for the correct goal location may be encoded in terms of trajectory representations, with the 'go left' plan encoded in the unique left-trajectory firing patterns occurring in the seconds leading up to the behavioral choice.

The RSC was remarkably sensitive to navigational goals, both in terms of highly specific responses at the reward locations, which were present during the earliest stages of learning (**Figure 23**), and in the form of simulations of upcoming reward locations as the rats approached the choice point. These results suggest an RSC role in memory and planning, which is consistent with studies of the default mode network. This network, which includes the RSC, prefrontal cortex, posterior parietal cortex, and hippocampus, mediates the constructive memory processes that underlie both episodic memory and the ability to imagine future events and situations (Buckner & Carroll, 2007; Schacter et al., 2012). Notably, many of these regions are also involved in route planning in humans (Brown et al., 2016; Ino et al., 2007; Kaplan et al., 2017; Takahashi et al., 1997), and evidence consistent with the representation of future routes and goal

locations has been reported in several of the same regions in rats (Ito et al., 2015; Johnson & Redish, 2007; Pfeiffer & Foster, 2013; Steiner & Redish, 2012). In our data, simulations of both goal locations emerged as the rats learned, and they only became selective for the correct location after the rat reached asymptotic levels of performance. Together with our finding that RSC lesions specifically impaired task performance at asymptote, these data suggest that future reward simulations in the RSC contribute to route planning, but only after other RSC spatial representations have become sufficiently stable.

The factors that drive the development of RSC spatial representations are not known. However, the functional similarities and anatomical connectivity between the hippocampus and RSC (Miller et al., 2014) suggest that the slow emergence of RSC representations may reflect consolidation of information from the hippocampus. Consistent with this idea, contextual fear memories depend on the hippocampus early after learning but later become more reliant on the RSC (Katche et al., 2010; Katche et al., 2013). In one particularly striking example, optogenetic reactivation of an RSC context representation was sufficient to evoke a contextual fear memory, even when the hippocampus was inactivated (Cowansage et al., 2014). Systems consolidation theory holds that spatial and episodic memories, which initially depend on the hippocampus, are eventually transferred to distributed cortical representations (McClelland, McNaughton, & O'Reilly, 1995; Squire & Alvarez, 1995) and several of our observations are consistent with that idea. RSC lesions had no discernable effect on the early stages of learning and only impaired performance after the task was well learned. Moreover, the performance deficit, though modest, was remarkably well correlated with the amount of tissue damage and spatial representations were spread across a wide swath (~5mm) of cortex with no obvious functional segregation, consistent with a widely distributed representation. The cortical memory representations that

support spatial navigation likely extend beyond the RSC, to include other midline cortical regions such as the anterior cingulate (Remondes & Wilson, 2015) and prefrontal cortex(Ito et al., 2015). Finally, the RSC representations themselves took a highly distributed form, with minimal information encoded by individual neurons and complex representations emerging only at the population level (McClelland et al., 1995).

Methods

Subjects

Subjects were 32 adult male Long Evans rats (Charles River Laboratories, Wilmington, MA) weighing 250g-300g upon arrival. Twelve rats were used in the electrophysiology study and 20 rats were used in the lesion study. Of the 10 rats that received RSC lesions, 2 were excluded from the analysis due to hippocampal damage, and 1 was excluded because the RSC damage was unilateral. Rats were placed on a 12hr/12hr light/dark cycle with lights on at 8am and allowed to acclimate to the vivarium for at least one week prior to surgery. After surgery, rats were placed on food restriction until they reached 80-85% of their free-feeding weight. Water was always available ad libitum. All procedures complied with the guidelines of the Cornell University Animal Care and Use Committee.

Surgery

Lesions. Twenty rats were anethetized with isoflorane gas (1-5% in oxygen) and placed in a Kopf stereotaxic apparatus. The skin was retracted and holes were drilled through the skull above each of the injection sites. Ten rats received bilateral neurotoxic (N-methyl-D-aspartate [NMDA], 10µg/ml) lesions of the RSC. NMDA was injected by hand in volumes of 0.20 – 0.35 µl using a custom-made glass injection canula (100µm diameter) attached to a Hamilton Syringe by sterile plastic tubing. The stereotaxic coordinates and injection volumes were:

1. 0.35µL at -2.2 (AP), ±0.5 (ML), -3.0 (DV)
2. 0.35µL at -3.9 (AP), ±0.5 (ML), -3.0 (DV)
3. 0.20µL at -5.5 (AP), ±0.5 (ML), -3.5 (DV)

4. 0.35µL at -5.5 (AP), ±1.0 (ML), -2.8 (DV)

5. 0.35µL at -6.7 (AP), ±1.1 (ML), -2.8 (DV)

6. 0.30µL at -8.0 (AP), ±1.3 (ML), -2.8 (DV)

Coordinates were taken from Bregma (AP), the midline (ML), and the surface of the skull (DV), respectively. The injection cannula was left in place 1 min before and 5 min after each infusion. An additional ten rats received sham lesions of the RSC consisting of lowering the injection cannula into the brain but not injecting NMDA.

Electrophysiology. Twelve rats had a custom-built electrode microdrive implanted, which contained 20 moveable tetrodes (16 recording tetrodes and 4 reference tetrodes) made from twisting four 17µm platinum/iridium (90%/10%) wires, platinum plated to an impedance of 100-500kΩ, and arranged in two 10-tetrode linear arrays (one in each hemisphere) that spanned approximately 5mm along the rostrocaudal axis of the brain. Tetrodes were stereotaxically positioned bilaterally just beneath the cortical surface (2-7mm posterior to Bregma, ±1.5mm lateral) with the tetrodes angled 30 degrees toward the midline. Rats were given 7 days to recover from surgery prior to lowering the tetrodes into the RSC (35-70µm daily) over the course of several days until a depth of at least 1mm was reached in order to ensure that the tetrodes were in the Rgb (discussed below).

Apparatus

The behavioral apparatus was a black PVC continuous T-maze (120 cm long stem x 100 cm wide x 68 cm above the floor) with metal reward cups embedded in the ends of the arms. Chocolate milk (0.2 ml, Nestle's Nesquik) could be delivered to the reward cups via an elevated

reservoir controlled by solenoid valves activated by foot-pedal switches. The maze was located in the center of a circular arena enclosed by black curtains with visual cues of various shapes, sizes, and colors. The room was illuminated by a ring of LED lights around the edge of the ceiling. A continuous background masking noise was played from a speaker located directly above the maze.

Behavioral Training Procedures

Prior to training, rats were acclimated to the maze and chocolate milk rewards with daily periods of free exploration on the maze until rats consumed 20 rewards within the first 10min of an acclimation session (mean = 4.5 acclimation days). After acclimating to the maze, rats were trained on a continuous spatial alternation task in which the rats were rewarded only if they approached the opposite (left or right) reward location from the previous trial. Both cups were baited on the first trial. Entries into the same arm as the previous trial were scored as an error and were not rewarded. Unlike some previous studies (Wood et al., 2000), rats were not shaped with trials where the incorrect choices were prevented by blocking access. Instead, rats were gently ushered back if they left the continuous alternation route. Rats were not allowed to correct their errors. Rats were given 40 trials/day until they achieved a criterion of 90% correct on two consecutive days. After achieving this criterion, rats were given up to 10 additional training sessions in order to record neuronal activity during asymptotic performance (i.e. overtraining).

Recordings

Neuronal spike data and video data were collected throughout learning (Digital Cheetah Data Acquisition System, Neuralynx, Inc. Bozeman, MT), filtered at 600Hz and 6kHz, digitized and

stored to disc along with timestamps for offline sorting (SpikeSort3D, Nueralynx, Inc.). The rat's position was monitored by digitized video of an LED array attached to the rat's head. The time of reward receipt was measured with a grounding circuit that detected oral contact with the chocolate milk reward.

Histology

After completion of the experiment, rats were transcardially perfused with 4% paraformaldehyde in phosphate buffered saline. Brains were removed and stored for at least 24hrs in 4% paraformaldehyde before being transferred to a 30% sucrose solution for storage until slicing. Coronal sections (40μm) were stained with 0.5% cresyl violet for visualization of tissue damage (in the case of NMDA lesions) or tetrode tracks (for electrophysiology recording implants). Tissue damage was quantified by laying a grid (250μm to-scale grid spacing) over an enlarged image of the stained tissue and dividing the number of grid intersections located over damaged RSC areas by the number of intersections located over the entire RSC. Boundaries of the RSC were determined in accordance with The Rat Brain in Stereotaxic Coordinates(Paxinos & Watson, 1998). Tetrode positions were identified using depth records noted during tetrode lowering and tracks observed in the stained tissue (**Figure 21**). Neuronal records from tetrodes located outside of the RSC were excluded from the data set. As in our previous work (Vedder et al., 2016), our recordings targeted the Rgb, although small numbers of neurons from the Rga or the RSA were also included. There were no conspicuous differences in the firing properties of neurons recorded in different subregions, different hemispheres, or at different AP coordinates.

Data analysis

Spatial coding. Because spatial responses have been extensively described at the single unit level (e.g., place fields), but RSC representations were more clearly seen at the population level, it was useful to assess neural responses at both levels (using previously established procedures for single unit analyses) and compare the results.

Spatial firing rate maps for each neuron were constructed by summing the total number of spikes that occurred in a spatial bin (2.5 x 2.5cm) and dividing by the amount of time spent in that bin. Bins visited on only one trial were excluded. Firing rate maps were smoothed by convolution with a 5 X 5 pixel Gaussian kernel. Neurons were classified as place cells if they had at least one place field, defined as a contiguous field of elevated firing that (1) was at least $50cm^2$ in size, but no larger than half the visited area of the maze, (2) had a within-field firing rate at least twice the firing rate outside the field, and (3) fired during more than 50% of the passes through the field. Both correct and error trials were included. Spatial information content was calculated (Ito et al., 2015) from firing rate maps computed as above, but smoothed using an adaptive method designed to optimize the trade-off between blurring error and sampling error (Skaggs, McNaughton, Wilson, & Barnes, 1996).

Bayesian decoding was used to predict the current position of the rat on the maze given the spiking activity of simultaneously recorded RSC populations and a uniform prior (Zhang et al., 1998) (**Figure 19E**). This analysis was restricted to recording sessions with at least 8 RSC neurons (34 sessions). Decoding was performed iteratively using a trial-based procedure whereby spike counts during time bins (200ms taken every 50ms) from one trial were used as the test sample, while the bins from all other trials were used as the training sample. The training sample was used to calculate firing rate maps for every neuron over a 50x50 grid overlaying the maze (mean of 352 visited pixels with each pixel approximately 2.5 X 2.5 cm). Probability

distributions of spike counts for each neuron and pixel were computed based on the mean spike counts and assuming a Poisson distribution. For each time bin in the test sample, the probability of the rat being in any particular pixel was calculated by multiplying, across neurons, the conditional probabilities of observing those spike counts if the rat occupied that pixel. The highest probability pixel was taken as the decoded position of the rat on the maze, and was considered an instance of correct decoding if it was within 4.5cm of the rat's actual position (i.e. within a circle with a diameter of approximately half the body length of a rat, ~9cm). Decoding accuracy was compared to a distribution of chance accuracies obtained by shuffling 10,000 times the spike counts of each neuron independently among the time bins for each recording session in that learning stage. The observed accuracy was considered significant if it was greater than 97.5% of the shuffle outcomes.

Correlation matrices were created to quantify the selectivity and reliability of RSC spatial firing throughout learning (**Figure 19G-I**). A single lap around the maze, beginning after the stem on a go-right trial and ending in the start area after a go-left trial, was divided into 170 spatial bins by partitioning observed paths through each section of the maze (**Figure 19G**). Standardized mean firing rate vectors were then calculated for each spatial bin independently for the first and second half of each session. To maximize comparability between learning stages, which had varying numbers of recorded neurons and systematic differences in behavior (i.e. more variable behavior was seen during early learning stages), the firing rate vectors were assembled from the first 50 neurons recorded during that stage and the analysis was limited to trials where the rat made typical passes through the maze section. Typical passes were defined as path lengths through a maze section that were shorter than 50% of all observed path lengths through that section (see path overlays in **Figure 19H**). Sessions with fewer than five typical

passes for each trial type (left and right) and session half were excluded. Separate correlation matrices were then generated for each learning stage (early, middle, late, and overtraining). Each row of a correlation matrix corresponds to the correlation of the ensemble rate vector for one spatial bin during the first half of the session with the ensemble rate vectors for every spatial bin during the second half of the session. Total spatial coding error was computed by finding all spatial bins in the first half of the session where the maximum correlation was not with itself in the second half (i.e. maximum correlations that were off the diagonal), and then summing the distances from the second-half bins back to the diagonal. To be conservative, the maze was treated as circular for the purpose of computing distance, and the shorter of the two distances (forward or backward) between the reconstructed and actual positions was always used. The observed spatial coding was compared to a chance distribution of spatial coding errors computed by shuffling the first-half second-half neuron pairings. To then determine the statistical significance of differences between training stages we compared the observed differences (in terms of mean squared error, MSE) to a distribution of differences obtained by shuffling the 200 neurons (50 per learning session) randomly between the four stages 10,000 times and each time recalculating the total spatial coding errors for each stage and the MSE between them. The observed MSE was considered significant if it was greater than 97.5% of shuffle-generated MSEs.

Location-specific reward responses. Location specific reward responses of individual RSC neurons were identified by comparing firing rates occurring in the 500ms before and after reward onset (oral contact with the milk) on every correct trial. Neurons were classified as having location-specific reward responses if they showed a significant ($p < 0.05$) interaction in a two-

way ANOVA (go-left/go-right trial type X Before/After 500ms time windows) and a significant post hoc ($p < 0.05$ t-test) comparing the 500ms after left and right rewards(Smith et al., 2012). In order to avoid comparing learning stages with different numbers of correct trials, only the first 10 correct trials of each trial type from each session were included in the analysis. To determine, within a learning stage, whether we observed a greater proportion of neurons with location-specific reward responses than was expected by chance, we compared the observed proportion to 10,000 control proportions obtained by randomly assigning each of the 20 trials as either left or right. The observed proportion was considered significant if it was greater than 97.5% of the control proportions. To then determine whether the observed proportion of neurons with location-specific reward responses differed significantly between the four learning stages, we shuffled individual neurons between stages 10,000 times, and calculated the MSE of the proportions after each shuffle. The observed MSE was considered statistically significant if it was greater than 97.5% of the shuffled MSEs.

To assess the specificity of ensemble firing at the reward locations, we combined neurons across rats and sessions to form ensemble firing rate vectors for visits to the left and right reward locations, and then computed a specificity measure that quantified how similar each visit was to other reward visits at the same location (e.g., left vs. left) and to the opposite location (e.g., left vs. right). Only the first 13 correct left and the first 13 right trials from each learning session were included (lowest number in any subject). To do this comparison, we used an iterative procedure whereby we excluded one trial from the data set, calculated mean left and right firing rate vectors from the remaining trials, and then computed the standardized Euclidean distance between the excluded trial and the two means. Specificity was then computed as the difference between the two distances normalized by the total distance. A positive specificity was counted as

a correct classification (i.e. minimum distance classification). Classification accuracy for each stage was compared to a chance distribution calculated by shuffling trial type labels (left or right) 10,000 times, and the observed classification accuracy was considered significant if it was greater than 97.5% of the control classification accuracies. To determine whether the observed classification accuracies differed between the four learning stages, we shuffled individual neurons between stages 10,000 times, and then calculated the MSE of the four control classification accuracies after each shuffle. The observed MSE was considered statistically significant if it was greater than 97.5% of the shuffled MSEs. We confirmed these classification results by applying principal component analysis to the same ensemble rate vectors and then performing linear discriminant analysis on the first 10 principal components(Vedder et al., 2016).

Trial-type specific firing on the stem. Individual neurons exhibiting trial-type specific firing as the rat traversed the stem were identified by comparing firing between left and right trials in each of four equal sized stem sectors (see **Figure 25A**, **Figure 26A**) using a two-way, repeated-measures ANOVA (Wood et al., 2000). Analyses were restricted to correct trials with typical stem runs that did not involve pauses or deviations from smooth locomotion (stem run times that exceeded 1.24s were excluded, 13.77% of learning trials and 3.63% of overtraining trials, see **Figure 26C**). To avoid comparing learning stages with different numbers of correct trials, only the first 10 correct trials of each trial type were included from each session. The statistical significance of the observed proportion of neurons was determined by shuffling 10,000 times both the (1) firing rates on each trial between the four sectors and (2) whether a trial was considered go-left or go-right, while maintaining the original proportion of each type. The

observed proportion was considered significant if it was greater than 97.5% of shuffle proportions. To then determine whether the observed proportion of neurons with location-specific reward responses differed significantly between the four learning stages, we shuffled individual neurons between stages 10,000 times, and calculated the MSE of the proportions after each shuffle. The observed MSE was considered statistically significant if it was greater than 97.5% of the shuffled MSEs.

Trial-type analyses of ensemble activity occurring on the stem were performed in the same manner as the above analyses of ensemble reward responses, but with trial type as the classification category variable instead of location. Only typical stem runs were included. Only the first 13 correct trials of each trial type were included for comparisons over learning. The first 17 trials of each type were included for the overtraining-only analyses due to greater numbers of correct trials (**Figure 27F-H**). Trial-type specificity and trial-type classification were then computed as above for ensemble responses to the rewards. Classification accuracy for each stage was compared to a chance distribution calculated by shuffling trial type labels 10,000 times, and the observed classification accuracy was considered significant if it was greater than 97.5% of the control classification accuracies. To determine whether the observed classification accuracies differed between the four learning stages, we shuffled individual neurons between stages 10,000 times, and then calculated the MSE of the four control classification accuracies after each shuffle. The observed MSE was considered statistically significant if it was greater than 97.5% of the shuffled MSEs. As above, we confirmed these classification results by applying principal component analysis to the same ensemble rate vectors and then performing linear discriminant analysis on the first 10 principal components.

Decoding to Future Reward Areas. Analyses of reward representations during stem traversals were similar to the above Bayesian analysis of spatial coding except that we sought to determine the degree to which the two reward locations were represented in the population activity rather than the rat's actual current position on the stem (**Figure 27A-C**). The analysis only included correct trials and the test sample was restricted to time bins as the rat traversed the stem. For each trial, we calculated the decoded probability (i.e. decoding) that the rat was located in the reward areas. Reward areas included both the reward locations and the portion of the approach arms after the choice point (see boxes in **Figure 27B**). Most reward area decoding was at or near the reward locations, but substantial decoding was also seen along the arms. To determine whether the reward areas were overrepresented relative to other non-stem areas, we normalized the amount of decoding to the reward areas by their relative size (proportion decoding divided by proportion of pixels) and compared the observed value to chance (i.e. a uniform distribution, proportion of decoded probability is equal to proportion of total pixels; dotted line in Figure 4d). The statistical significance of each stage mean was calculated by comparing the observed distribution of session means to a value of 1.0 using a Bonferroni-corrected one-sample t-test. Representations of the two reward areas (left and right) were then compared to each other to determine whether the rat preferentially represented the reward area that it was about to approach (**Figure 27E**). The difference between the representations (decoded probabilities) of the correct and incorrect reward areas were computed and then standardized by their sum (correct minus incorrect divided by the total). Positive values indicate a greater representation of the correct reward area, while negative values indicate a greater representation of the incorrect reward area. The statistical significance of each stage mean was calculated by comparing the observed distribution of session means to zero as above.

Acknowledgements

I am particularly grateful for the contributions of William Mau, who completed his undergraduate thesis while assisting on this project. I also thank Sarah Parauda, Keunhyung Yu, Alexandra Tse, and Hei Jun Li for assistance with animal training and electrode hyperdrive fabrication.

References

Alexander, A. S., & Nitz, D. A. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nat Neurosci, 18*(8), 1143-1151. doi: 10.1038/nn.4058

Alexander, A. S., & Nitz, D. A. (2017). Spatially Periodic Activation Patterns of Retrosplenial Cortex Encode Route Sub-spaces and Distance Traveled. *Current Biology*. doi: 10.1016/j.cub.2017.04.036

Auger, S. D., & Maguire, E. A. (2013). Assessing the mechanism of response in the retrosplenial cortex of good and poor navigators. *Cortex, 49*(10), 2904-2913. doi: 10.1016/j.cortex.2013.08.002

Auger, S. D., Mullally, S. L., & Maguire, E. A. (2012). Retrosplenial cortex codes for permanent landmarks. *PLoS One, 7*(8), e43620. doi: 10.1371/journal.pone.0043620

Auger, S. D., Zeidman, P., & Maguire, E. A. (2015). A central role for the retrosplenial cortex in de novo environmental learning. *eLife, 4*, e09031. doi: 10.7554/eLife.09031

Brown, T. I., Carr, V. A., LaRocque, K. F., Favila, S. E., Gordon, A. M., Bowles, B., . . . Wagner, A. D. (2016). Prospective representation of navigational goals in the human hippocampus. *Science, 352*(6291), 1323-1326. doi: 10.1126/science.aaf0784

Buckner, R. L., & Carroll, D. C. (2007). Self-projection and the brain. *Trends Cogn Sci, 11*(2), 49-57. doi: 10.1016/j.tics.2006.11.004

Burke, S. N., Maurer, A. P., Nematollahi, S., Uprety, A. R., Wallace, J. L., & Barnes, C. A. (2011). The influence of objects on place field expression and size in distal hippocampal CA1. *Hippocampus, 21*(7), 783-801. doi: 10.1002/hipo.20929

Bussey, T. J., Muir, J. L., Everitt, B. J., & Robbins, T. W. (1996). Dissociable effects of anterior and posterior cingulate cortex lesions on the acquisition of a conditional visual discrimination: facilitation of early learning vs. impairment of late learning. *Behav Brain Res, 82*(1), 45-56.

Buzsaki, G., & Moser, E. I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat Neurosci, 16*(2), 130-138. doi: 10.1038/nn.3304

Chen, L. L., Lin, L. H., Green, E. J., Barnes, C. A., & McNaughton, B. L. (1994). Head-direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. *Exp Brain Res, 101*(1), 8-23.

Cho, J., & Sharp, P. E. (2001). Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behav Neurosci, 115*(1), 3-25.

Chun, M. M. (2000). Contextual cueing of visual attention. *Trends Cogn Sci, 4*(5), 170-178.

Corcoran, K. A., Donnan, M. D., Tronson, N. C., Guzman, Y. F., Gao, C., Jovasevic, V., . . . Radulovic, J. (2011). NMDA receptors in retrosplenial cortex are necessary for retrieval of recent and remote context fear memory. *J Neurosci, 31*(32), 11655-11659. doi: 10.1523/jneurosci.2107-11.2011

Cowansage, K. K., Shuman, T., Dillingham, B. C., Chang, A., Golshani, P., & Mayford, M. (2014). Direct reactivation of a coherent neocortical memory of context. *Neuron, 84*(2), 432-441. doi: 10.1016/j.neuron.2014.09.022

Czajkowski, R., Jayaprakash, B., Wiltgen, B., Rogerson, T., Guzman-Karlsson, M. C., Barth, A. L., . . . Silva, A. J. (2014). Encoding and storage of spatial information in the retrosplenial cortex. *Proc Natl Acad Sci U S A, 111*(23), 8661-8666. doi: 10.1073/pnas.1313222111

Day, M., Langston, R., & Morris, R. G. (2003). Glutamate-receptor-mediated encoding and retrieval of paired-associate learning. *Nature, 424*(6945), 205-209. doi: 10.1038/nature01769

Frank, L. M., Brown, E. N., & Wilson, M. (2000). Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron, 27*(1), 169-178.

Gabriel, M. (1993). Discriminative avoidance learning: a model system *Neurobiology of cingulate cortex and limbic thalamus* (pp. 478-523). Boston: Birkhäuser.

Ino, T., Doi, T., Hirose, S., Kimura, T., Ito, J., & Fukuyama, H. (2007). Directional disorientation following left retrosplenial hemorrhage: a case report with fMRI studies. *Cortex, 43*(2), 248-254.

Ito, H. T., Zhang, S. J., Witter, M. P., Moser, E. I., & Moser, M. B. (2015). A prefrontal-thalamo-hippocampal circuit for goal-directed spatial navigation. *Nature, 522*(7554), 50-55. doi: 10.1038/nature14396

Jacob, P. Y., Casali, G., Spieser, L., Page, H., Overington, D., & Jeffery, K. (2017). An independent, landmark-dominated head-direction signal in dysgranular retrosplenial cortex. *Nat Neurosci, 20*(2), 173-175. doi: 10.1038/nn.4465

Johnson, A., & Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *J Neurosci, 27*(45), 12176-12189. doi: 10.1523/jneurosci.3761-07.2007

Kaplan, R., King, J., Koster, R., Penny, W. D., Burgess, N., & Friston, K. J. (2017). The Neural Representation of Prospective Choice during Spatial Planning and Decisions. *15*(1), e1002588. doi: 10.1371/journal.pbio.1002588

Katche, C., Bekinschtein, P., Slipczuk, L., Goldin, A., Izquierdo, I. A., Cammarota, M., & Medina, J. H. (2010). Delayed wave of c-Fos expression in the dorsal hippocampus involved specifically in persistence of long-term memory storage. *Proc Natl Acad Sci U S A, 107*(1), 349-354. doi: 10.1073/pnas.0912931107

Katche, C., Dorman, G., Gonzalez, C., Kramar, C. P., Slipczuk, L., Rossato, J. I., . . . Medina, J. H. (2013). On the role of retrosplenial cortex in long-lasting memory storage. *Hippocampus, 23*(4), 295-302. doi: 10.1002/hipo.22092

Maguire, E. A., Nannery, R., & Spiers, H. J. (2006). Navigation around London by a taxi driver with bilateral hippocampal lesions. *Brain, 129*(Pt 11), 2894-2907. doi: 10.1093/brain/awl286

Marchette, S. A., Vass, L. K., Ryan, J., & Epstein, R. A. (2014). Anchoring the neural compass: coding of local spatial reference frames in human medial parietal lobe. *Nat Neurosci, 17*(11), 1598-1606. doi: 10.1038/nn.3834

Maviel, T., Durkin, T. P., Menzaghi, F., & Bontempi, B. (2004). Sites of neocortical reorganization critical for remote spatial memory. *Science, 305*(5680), 96-99. doi: 10.1126/science.1098180

McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev, 102*(3), 419-457.

Miller, A. M. P., Vedder, L. C., Law, L. M., & Smith, D. M. (2014). Cues, context, and long-term memory: the role of the retrosplenial cortex in spatial cognition. *Frontiers in Human Neuroscience, 8*. doi: 10.3389/fnhum.2014.00586

Morgan, L. K., Macevoy, S. P., Aguirre, G. K., & Epstein, R. A. (2011). Distances between real-world locations are represented in the human hippocampus. *J Neurosci, 31*(4), 1238-1245. doi: 10.1523/jneurosci.4667-10.2011

Nelson, A. J., Powell, A. L., Holmes, J. D., Vann, S. D., & Aggleton, J. P. (2015). What does spatial alternation tell us about retrosplenial cortex function? *Front Behav Neurosci, 9*, 126. doi: 10.3389/fnbeh.2015.00126

Paxinos, G., & Watson, C. (1998). *The rat brain in stereotaxic coordinates*.

Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature, 497*(7447), 74-79. doi: 10.1038/nature12112

Pothuizen, H. H., Davies, M., Aggleton, J. P., & Vann, S. D. (2010). Effects of selective granular retrosplenial cortex lesions on spatial working memory in rats. *Behav Brain Res, 208*(2), 566-575. doi: 10.1016/j.bbr.2010.01.001

Remondes, M., & Wilson, M. A. (2015). Slow-gamma Rhythms Coordinate Cingulate Cortical Responses to Hippocampal Sharp-Wave Ripples during Wakefulness. *Cell Rep, 13*(7), 1327-1335. doi: 10.1016/j.celrep.2015.10.005

Schacter, D. L., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., & Szpunar, K. K. (2012). The future of memory: remembering, imagining, and the brain. *Neuron, 76*(4), 677-694. doi: 10.1016/j.neuron.2012.11.001

Shine, J. P., Valdes-Herrera, J. P., Hegarty, M., & Wolbers, T. (2016). The Human Retrosplenial Cortex and Thalamus Code Head Direction in a Global Reference Frame. *J Neurosci, 36*(24), 6371-6381. doi: 10.1523/jneurosci.1268-15.2016

Skaggs, W. E., McNaughton, B. L., Wilson, M. A., & Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus, 6*(2), 149-172. doi: 10.1002/(SICI)1098-1063(1996)6:2<149::AID-HIPO6>3.0.CO;2-K

Smith, D. M., Barredo, J., & Mizumori, S. J. (2012). Complimentary roles of the hippocampus and retrosplenial cortex in behavioral context discrimination. *Hippocampus, 22*(5), 1121-1133. doi: 10.1002/hipo.20958

Squire, L. R., & Alvarez, P. (1995). Retrograde amnesia and memory consolidation: a neurobiological perspective. *Curr Opin Neurobiol, 5*(2), 169-177.

Steiner, A. P., & Redish, A. D. (2012). The road not taken: neural correlates of decision making in orbitofrontal cortex. *Front Neurosci, 6*, 131. doi: 10.3389/fnins.2012.00131

Takahashi, N., Kawamura, M., Shiota, J., Kasahata, N., & Hirayama, K. (1997). Pure topographic disorientation due to right retrosplenial lesion. *Neurology, 49*(2), 464-469.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychol Rev, 55*(4), 189-208.

Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., . . . Morris, R. G. (2007). Schemas and memory consolidation. *Science, 316*(5821), 76-82. doi: 10.1126/science.1135935

Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., . . . Morris, R. G. (2011). Schema-dependent gene activation and memory encoding in neocortex. *Science, 333*(6044), 891-895. doi: 10.1126/science.1205274

van Groen, T., & Wyss, J. M. (1992). Connections of the retrosplenial dysgranular cortex in the rat. *J Comp Neurol, 315*(2), 200-216. doi: 10.1002/cne.903150207

van Groen, T., & Wyss, J. M. (2003). Connections of the retrosplenial granular b cortex in the rat. *J Comp Neurol, 463*(3), 249-263. doi: 10.1002/cne.10757

Vann, S. D., Aggleton, J. P., & Maguire, E. A. (2009). What does the retrosplenial cortex do? *Nat Rev Neurosci, 10*(11), 792-802. doi: 10.1038/nrn2733

Vedder, L. C., Miller, A. M., Harrison, M. B., & Smith, D. M. (2016). Retrosplenial Cortical Neurons Encode Navigational Cues, Trajectories and Reward Locations During Goal Directed Navigation. *Cereb Cortex*. doi: 10.1093/cercor/bhw192

Wheeler, A. L., Teixeira, C. M., Wang, A. H., Xiong, X., Kovacevic, N., Lerch, J. P., . . . Frankland, P. W. (2013). Identification of a functional connectome for long-term fear memory in mice. *PLoS Comput Biol, 9*(1), e1002853. doi: 10.1371/journal.pcbi.1002853

Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science, 261*(5124), 1055-1058.

Winocur, G., Moscovitch, M., Fogel, S., Rosenbaum, R. S., & Sekeres, M. (2005). Preserved spatial memory after hippocampal lesions: effects of extensive experience in a complex environment. *Nat Neurosci, 8*(3), 273-275. doi: 10.1038/nn1401

Wolbers, T., & Buchel, C. (2005). Dissociable retrosplenial and hippocampal contributions to successful formation of survey representations. *J Neurosci, 25*(13), 3333-3340. doi: 10.1523/jneurosci.4705-04.2005

Wood, E. R., Dudchenko, P. A., Robitsek, R. J., & Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron, 27*(3), 623-633.

Zhang, K., Ginzburg, I., McNaughton, B. L., & Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J Neurophysiol, 79*(2), 1017-1044.

Figures

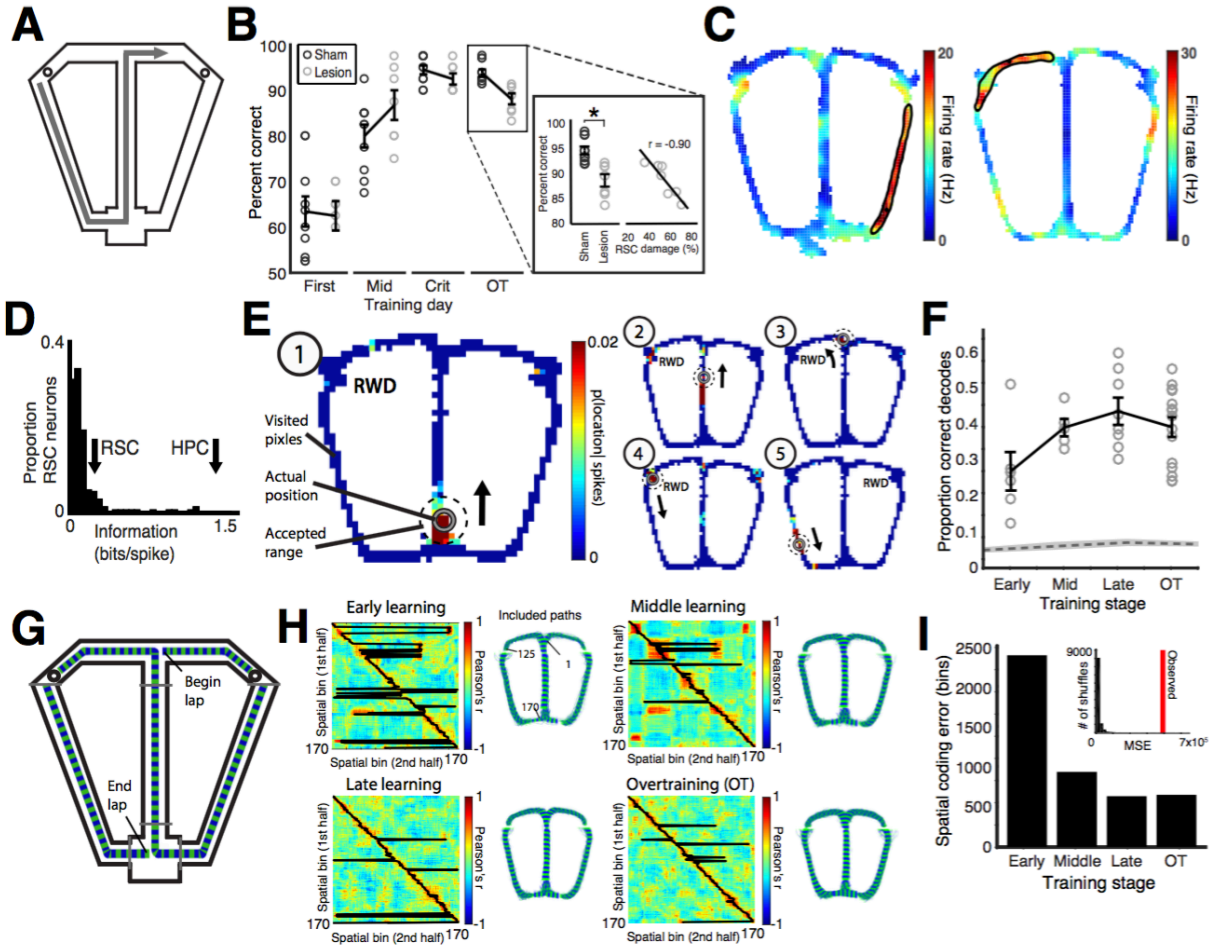

Figure 19 RSC ensembles developed a spatial representation of the maze with learning. (**A**) Schematic diagram of the continuous t-maze alternation task. After visiting one of the two reward locations (circles), rats returned to the stem and had to approach the opposite location for reward. (**B**) Permanent lesions of the RSC selectively impair spatial alternation performance after learning. Behavioral performance is plotted for the first (First), middle (Mid), and last (criterial, Crit) learning days, and asymptotic performance days (overtraining, OT). Performance for each control and lesion rat (left) is shown as open circles, with error bars showing +/- SEM. Inset shows OT performance, along with the correlation between performance and lesion size. (**C**) Firing rate maps of two representative RSC neurons with place fields. These neurons showed

increased firing in contiguous spatial areas (black outlines), but also showed considerable extra-field firing. (**D**) A histogram of information content scores for all RSC neurons recorded during overtraining. RSC arrow indicates the mean of RSC neurons meeting place field criteria. The mean hippocampal information content score for all pyramidal neurons from a similar task by Ito et al. (2015) is shown for comparison (HPC arrow). (**E**) A Bayesian decoder was applied to the activity of simultaneously recorded ensembles of RSC neurons to predict the rat's current location on the maze based on instantaneous firing. Five decoded instances from one trial are shown. Colored pixels in the shape of the maze indicate the probability of the rat being located in that pixel given the instantaneous spiking activity, with warmer colors corresponding to higher probabilities. The highest probability pixel was taken as the decoded location. The rat's actual location is shown by the grey circle, and the rat's current direction of travel is indicated by the black arrow. Decoded locations falling within the dashed circle were considered accurate. (**F**) RSC ensembles show improved encoding of spatial position with training. The decoder's success rates for individual sessions are plotted as open circles, while averages for each training stage are plotted with error bars showing +/- SEM. Decoding improved significantly with training and was always far more accurate than was expected by chance (gray area shows the center 95% of the shuffle distribution). (**G**) Improvements in spatial representation were not due to learning-related changes in behavior. A full lap around the maze was divided into 170 spatial bins along strictly defined behavioral trajectories, and correlations were computed between firing rate vectors from the first and second half of each session at all spatial bins. (**H**) Correlation matrices from early, middle, and late learning sessions, and from overtraining show correlations between activity occurring in each bin during the first half and second half of sessions. The black line connects the pixels of highest correlation between the two session halves at each spatial bin. Deviations

from the diagonal are indicative of spatial coding errors. Overlaid paths from all trials that contributed firing rate data are shown. (**I**) Total spatial coding error is plotted for each learning stage. Inset, the observed difference in coding errors between stages (red line) was far greater than the shuffled distribution.

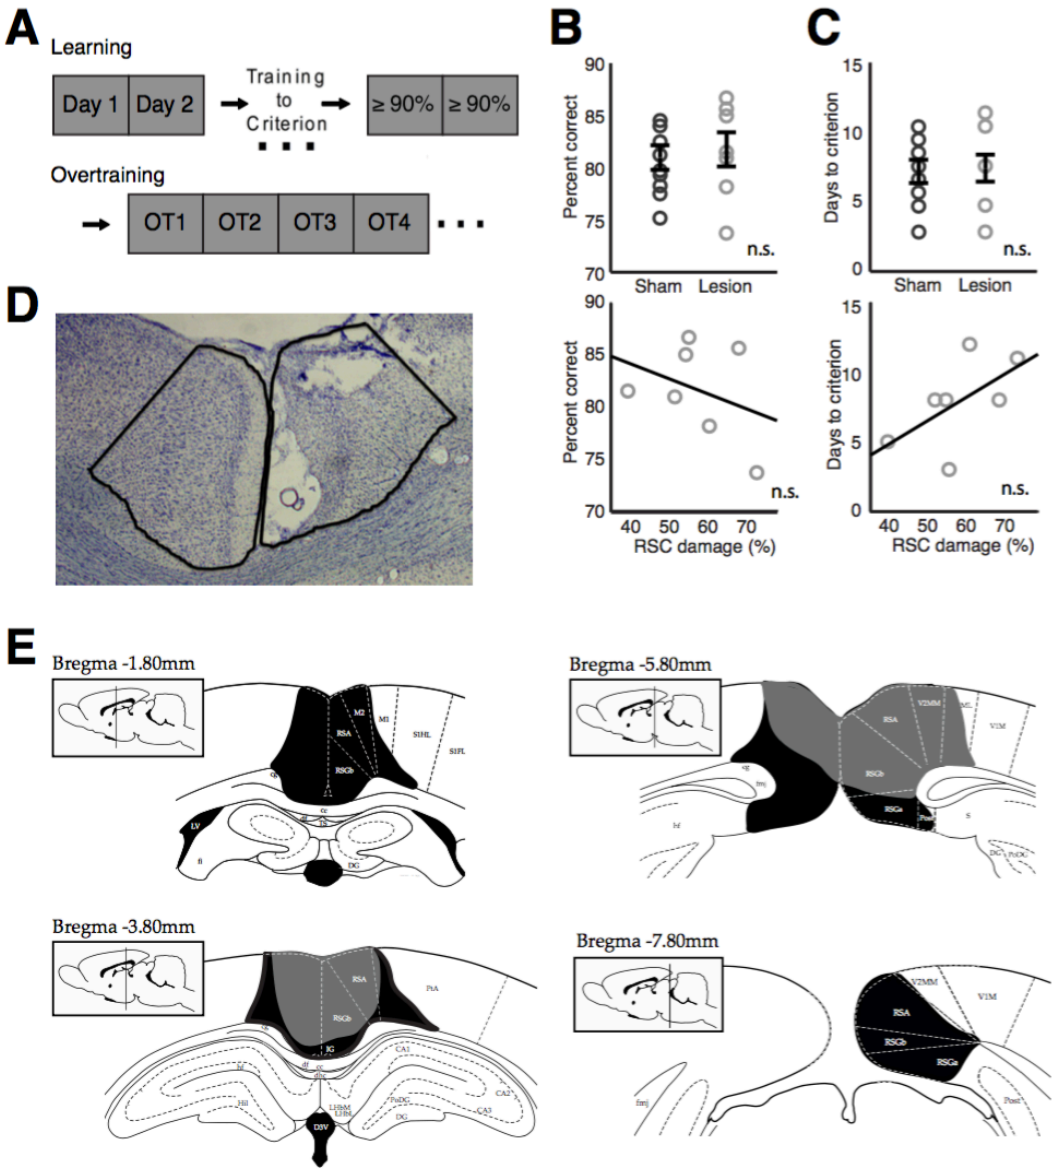

Figure 20 RSC lesions did not impair the acquisition of spatial alternation. (**A**) Diagram of the continuous spatial alternation training procedure. Rats were trained daily until they achieved a behavioral criterion of 90% on two consecutive days, and then were given at least four additional training days (i.e. overtraining, OT). (**B**) Rats with lesions of the retrosplenial cortex (RSC) showed similar alternation accuracy during learning (top, $t_{(15)} = 0.33$, $p = 0.75$), and there was not a significant correlation between alternation accuracy during learning and lesion size (bottom, $r = -0.36$, $p = 0.43$). Open circles correspond to individual rats' mean percent correct

over all learning days. Error bars show mean performance in each condition +/- SEM. (**C**) Rats with lesions of the RSC required a similar number of training days to reach criterion (top, $t_{(15)} = 0.19$, $p = 0.85$) and there was not a significant correlation between number of training days and lesion size (bottom, $r = 0.38$, $p = 0.40$). (**D**) A nissl stained coronal section of the RSC (black boundaries) from a rat with a unilateral lesion (not included in the analyses), illustrating the effects of NMDA excitotoxic lesions (right) and an intact RSC (left). (**E**) Damaged area observed in the rats with the largest (black) and smallest (gray) lesions are shown at 4 AP positions.

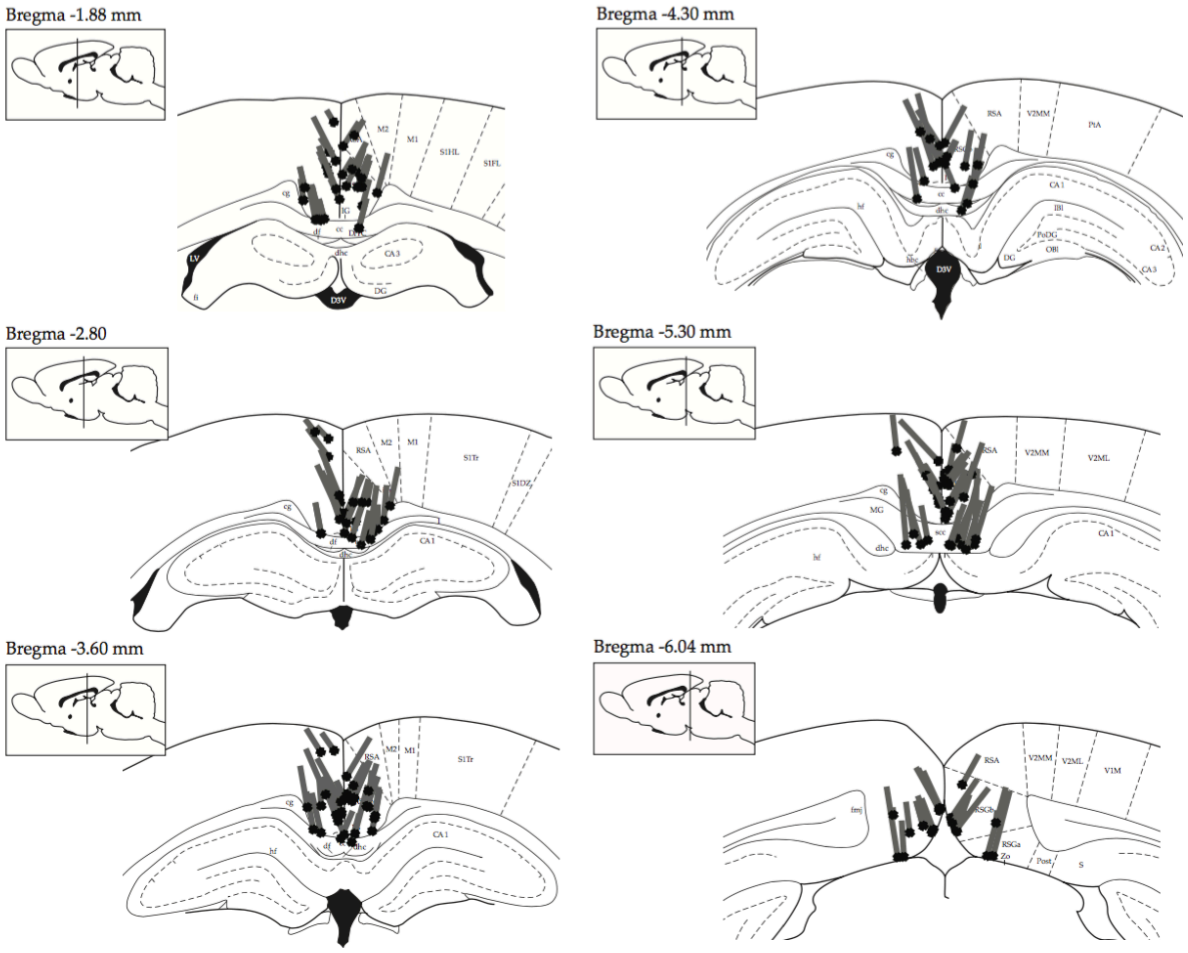

Figure 21 Estimated tetrode positions during recording (gray lines) and the corresponding end positions (black dots) are shown at 6 AP positions. End positions indicate the lowest point of the tetrode, and not necessarily the location of tetrode during the last recording session, which was often much earlier. Neurons were recorded primarily in Rgb, with smaller numbers of neurons recorded in the RSA and Rga. Any recordings from tetrodes located outside of the RSC were omitted from the analysis.

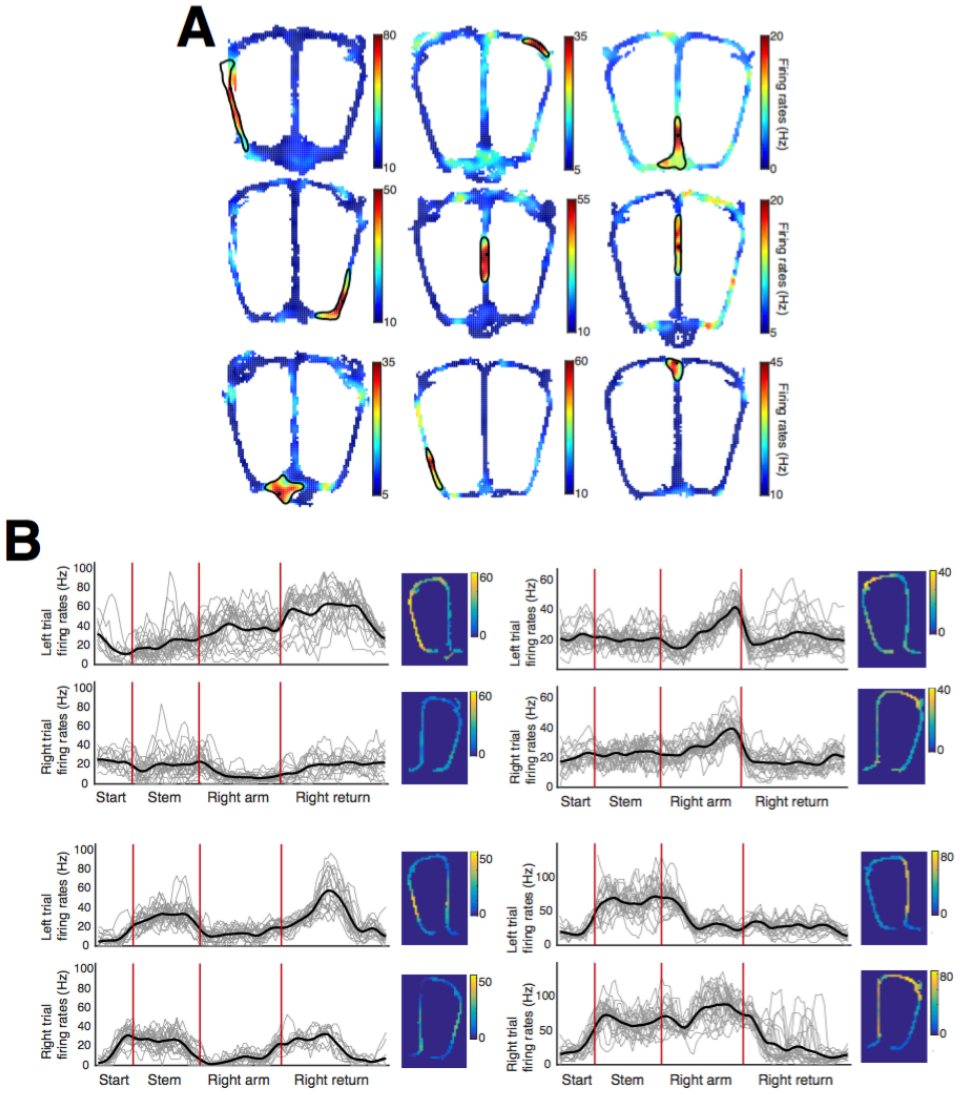

Figure 22 Spatial firing characteristics of individual RSC neurons. Some RSC neurons (~19%) could be classified as having place fields using previously published criteria (Smith, Barredo, & Mizumori, 2012). (**A**) Firing rate maps for nine example RSC neurons that met place field classification criteria are shown. Black outlines show the boundaries of the place fields. Note that the firing rate color keys do not start at zero, as most RSC neurons have high background firing rates outside of the defined place field and the resulting place fields are less specific than those of the hippocampus. RSC place fields in this data set were much larger (mean = 225cm^2, compared to 80cm^2 in the hippocampus (Burke et al., 2011)), had less contrast between infield

and out-of-field firing (infield-outfield firing ratios, mean = 3.25, compared to 50.34 in the hippocampus (Smith et al., 2012)), and RSC firing had reduced spatial information content (mean at asymptote = 0.122 bits/spike ± 0.010se; compared to 1.46 ± 0.09 in the HPC, from Ito et al., 2015). Despite the relatively poor spatial specificity of RSC neurons and high background firing rates, RSC neuronal firing was fairly reliable from one trial to the next. (**B**) The firing of four example neurons illustrating the trial by trial firing (grey lines) as rats traversed the maze on left and right trials, along with the average of all trials (black lines, see also (Alexander & Nitz, 2015)). Red vertical lines show maze section boundaries. For this figure, only typical passes through each section were included (as in **Figure 19H, insets**). Note that the firing rates often reliably distinguished between the maze sections and between left and right trial types. These firing patterns underlie the population code for spatial location illustrated in **Figure 19**. Firing rate maps separated into left and right trial types are shown to the right of each plot.

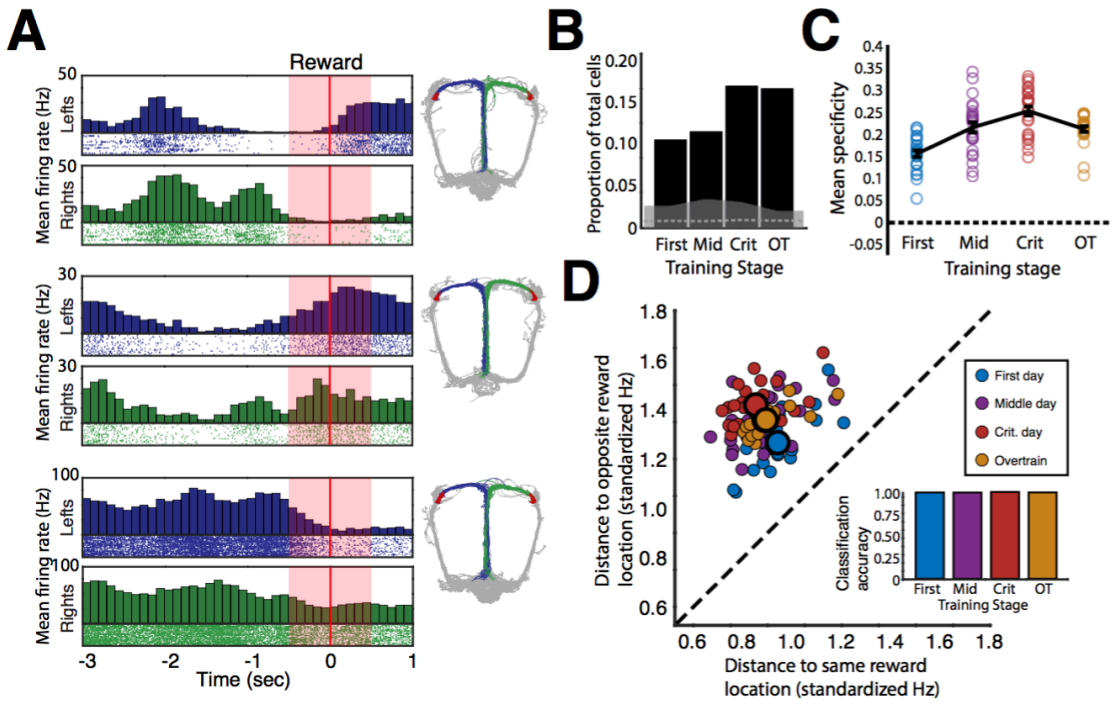

Figure 23 Neurons in the RSC uniquely encode each reward location. (**A**) Three examples of RSC neurons that encoded reward locations. Firing rate histograms and trial-by-trial raster plots of individual spikes show three seconds before to one second after the rat's first lick of the reward (red line) for left (blue) and right (green) trials. The shading indicates the 1s time window surrounding lick detection that was used to identify reward responses (see methods). Plots of the rats' paths on the maze are shown for all visited locations (gray lines), with positions during time windows used for the accompanying firing rate plots shown (blue and green overlays). The red portions of the paths correspond to the 1s time window that was used to identify reward responses (shaded area in the plots). (**B**) Proportions of retrosplenial neurons uniquely encoding each reward location are shown for each training stage (dotted line and gray shaded area indicate the mean and center 95% of the shuffle distribution). (**C**) Location specificity of RSC neural ensemble firing increased with training. Ensemble activity (combined across subjects and sessions, see methods) from individual reward visits is plotted in terms of mean specificity,

defined as the difference between the distances to the representations of the opposite and same reward location, standardized by the sum of both distances. Individual reward visits are plotted as open circles, while averages for each training stage are plotted with error bars showing +/- SEM. (**D**) Colored dots show ensemble activity from one trial (as with colored circles in **c**) plotted in terms of its distance from the mean of the same and opposite reward location representations. Points along the dotted line are equidistant to both reward locations, while points farther from the dotted line indicate stronger ensemble preferences for one reward location over the other. Large dots illustrate the mean for each learning stage. Note that ensemble activity reliably prefers the same reward location to the opposite location throughout learning, but also moves away from the unity line as learning progresses. The ability to classify trials (left or right reward locations) on the basis of ensemble firing was perfect throughout learning (inset).

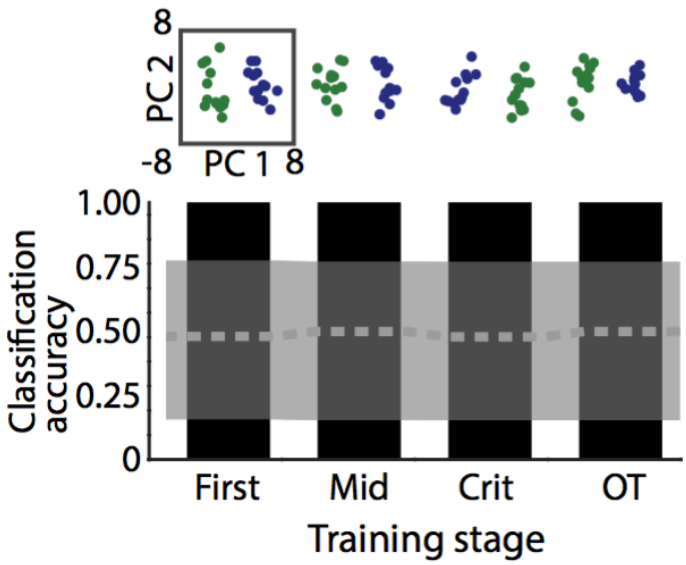

Figure 24 RSC ensembles clearly discriminated the two reward locations. Principal components analysis was used to reduce the dimensionality of RSC ensemble activity occurring during the 1s time window surrounding reward on left and right trials. The first two principal components are shown for each learning stage with green dots showing right trials and blue dots showing left trials. Linear discriminant analysis performed on the first 10 principal components revealed perfect discrimination between the left and right reward locations (see methods, chance accuracy is indicated by the dotted line and gray shading).

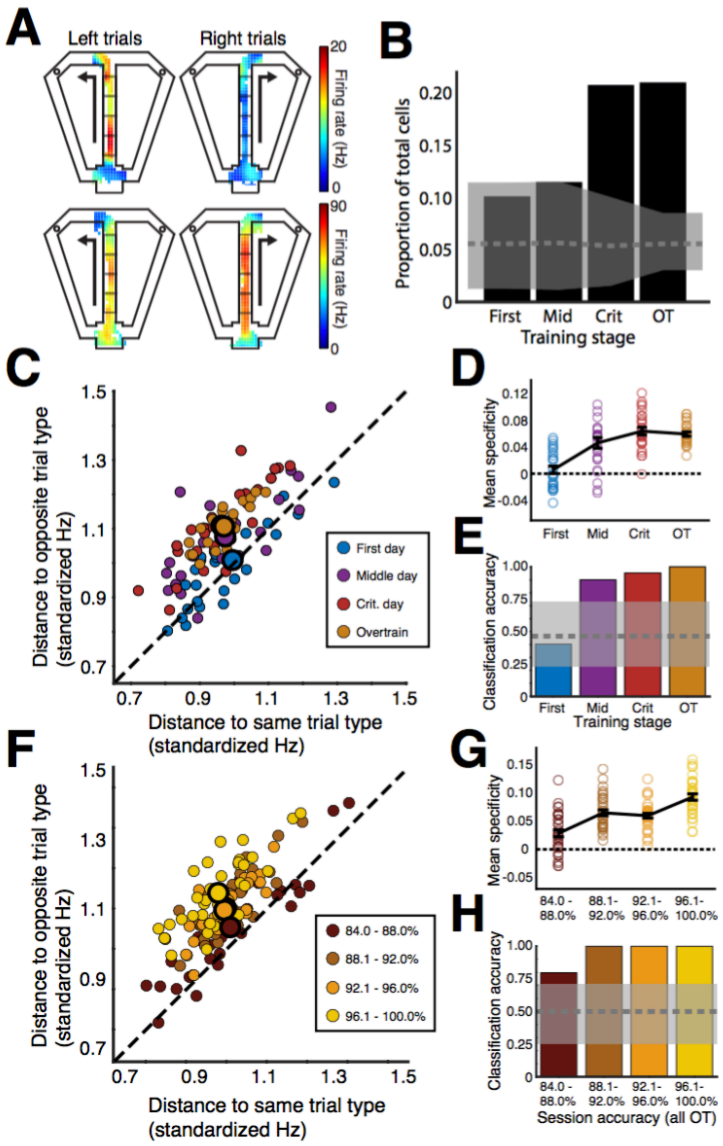

Figure 25 Neurons in the RSC develop trial-type specific firing on the stem. (**A**) Two examples of RSC neurons (rows) that fired differently on the stem depending on whether the rat was about to turn left or right (i.e. trial-type specific firing). Firing rate maps are shown for left and right turn trials, with the analyzed sectors of the stem indicated. (**B**) The proportion of RSC neurons showing trial-type specific firing is plotted for the first (First), middle (Mid), and last (criterion, Crit) learning day, and for all overtraining (OT) days. This proportion increased from chance levels (gray shading) on the first day to over 20% at criterion day and during overtraining. (**C**)

Trial-type specificity of RSC neural ensemble firing increased with training. Colored dots show ensemble activity from one trial (combined across subjects and sessions, see methods) plotted in terms of its distance from the mean of the same and opposite trial types. Points along the dotted line are equidistant to both trial types, indicating that the ensemble showed no preference for left or right trials, while points farther from the dotted line indicate stronger ensemble preferences for one trial type over the other. Large dots illustrate the mean for each learning stage. Note that ensemble activity diverges from the unity line as learning progresses. (**D**) Trial-type specificity of the RSC ensemble increased with training and was greater than chance by the middle training day. Individual trials (small dots from C) are plotted as open circles, while averages for each training stage are plotted with error bars showing +/- SEM. (**E**) The ability to classify trials (left or right) solely on the basis of ensemble firing patterns improved with learning, from chance (gray shading) on the first day of training, to perfect accuracy during overtraining. (**F-H**) The trial-type specificity of RSC ensemble firing was greater during sessions with better alternation performance. Plots are the same as C-E**,** except that all data were taken from overtraining sessions that were grouped according to behavioral performance (% correct choices for the session). Note that ensemble activity shows increased trial-type specificity and improved classification of left and right trials during superior behavioral performance.

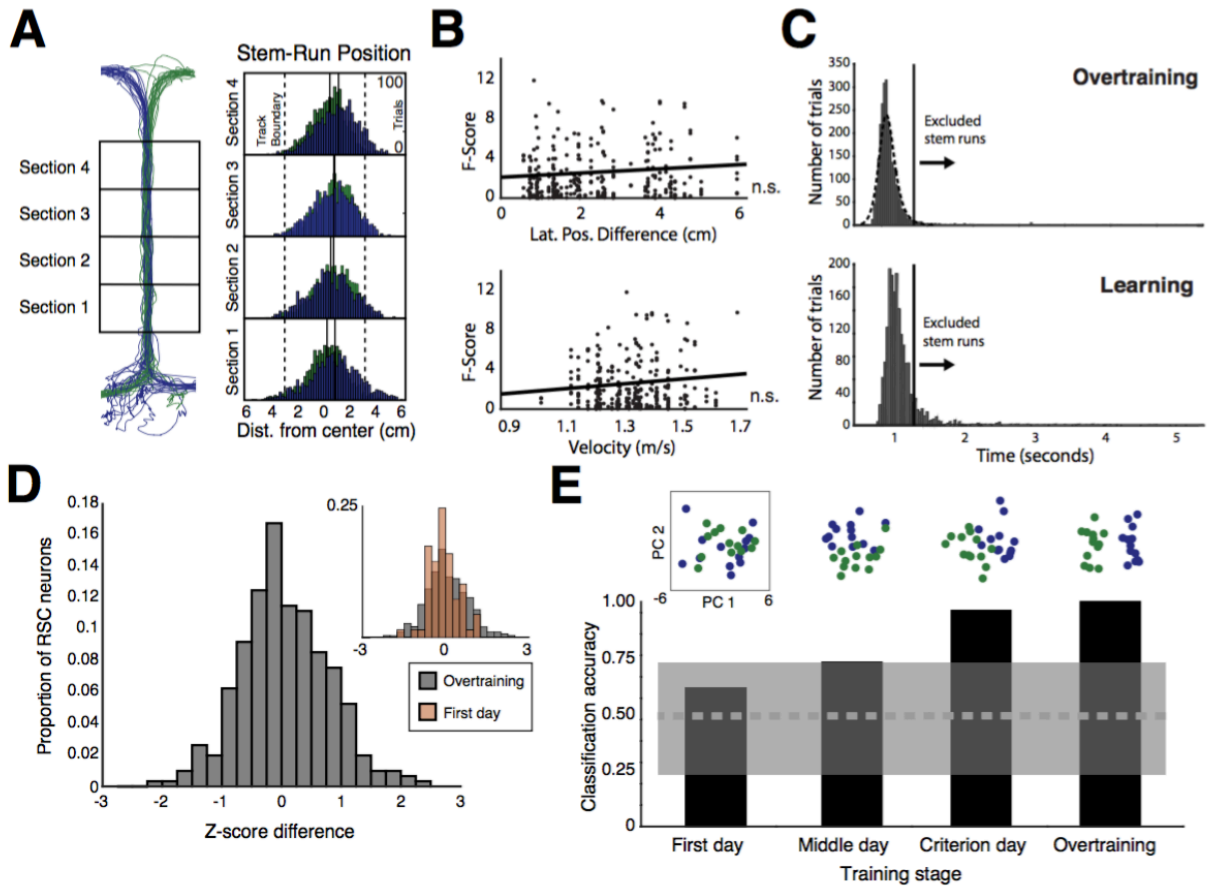

Figure 26 Behavior and trial-type specific firing on the stem of the maze. (**A**) Stem positions on left and right trial types were highly overlapping in each of the four stem sectors used to define trial-type specificity in individual neurons (see methods). All of the included stem runs from one overtraining session are plotted (left), with left trials shown in blue and right trials shown in green. The lateral position on the stem of all overtraining trials in the dataset with left and right trials plotted for each sector (right). Bars of the histogram correspond to the number of trials where a rat occupied that lateral position on the stem, the solid black lines indicate trial-type mean position, and the dotted lines indicate the left and right boundaries of the stem. Because we determined the rat's spatial position using lights on the head of the rat, the rat could be considered outside the stem boundaries if he tilted his head to the side. (**B**) Trial-type specific firing on the stem is not correlated with lateral position (top) or running speed (bottom). Similar

to the analysis for trial-type specificity in individual neurons (see methods), each neuron in the dataset was subjected to a 2 (trial type) X 4 (stem sector) repeated measures ANOVA and the F-score from either the main effect of trial type or the trial-type by stem sector interaction (whichever was greater) was used as a measure of trial-type specific firing. F-scores were not correlated with the lateral position difference between left and right trials (summed across all four sectors). F–scores were also not correlated with the average running speed on the stem during that session. (**C**) Time on the stem was used as a measure of behavioral typicality, with particularly high dwell times being indicative of atypical behavior. Histograms show the stem dwell times of every trial recorded during overtraining (top) and learning (bottom) sessions. To exclude atypical stem behavior, a log-logistic curve was fit to the distribution of all stem-run durations during overtraining (dotted line). The solid black lines indicate a cutoff at 2.5 standard deviations above the mean of the fitted curve, selected based on visual inspection of the overtraining data. This cut off (1.24s) was then used to exclude stem runs from both overtraining and learning sessions. (**D**) Histogram illustrating the distribution of trial-type specific firing in individual RSC neurons. For each neuron, we calculated the difference between the mean firing rate of all right trials and all left trials, and then divided this by the pooled standard deviation (i.e. the Z-score difference). Negative scores indicate higher firing on left trials, while positive scores indicate higher firing on right trials. Inset, for comparison, the distribution on the first day (orange) is overlaid on the overtraining distribution (gray). Note that the distribution becomes wider with training, as the trial-type specificity increases. (**E**) RSC populations developed the ability to discriminate between go-left and go-right trial types. Principal components analysis was used to reduce the dimensionality of RSC ensemble activity occurring on the stem during left and right trials (see methods). The first two principal components are shown for each

learning stage (dot plots). Note that the clusters of left (blue) and right (green) trials become more distinct over learning. Bottom, a linear discriminant analysis of the dimensionally reduced dataset (first ten principal components) revealed chance level (dotted line and gray shading indicate the mean and center 95% of a shuffled distribution) discrimination early in learning, but improved to perfect discrimination during overtraining.

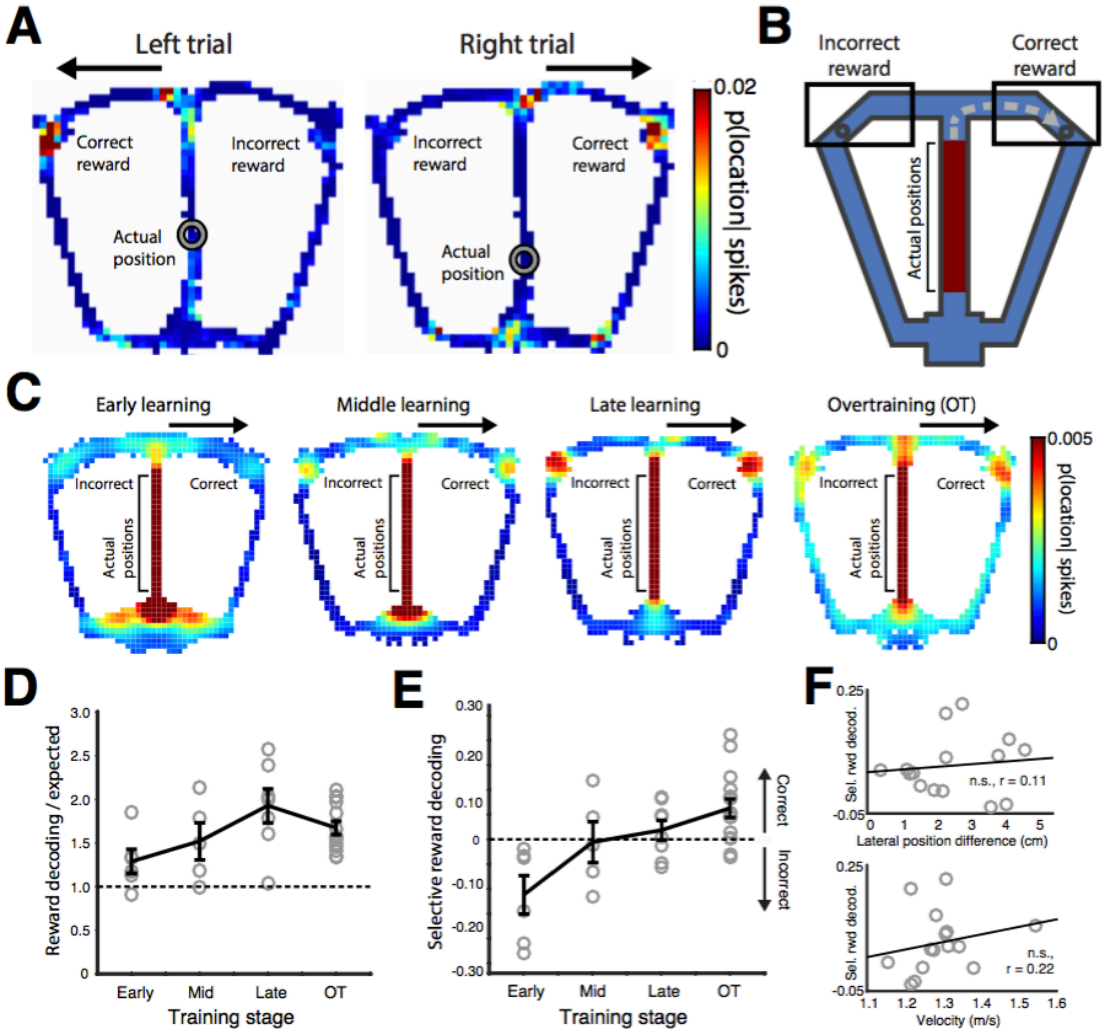

Figure 27 RSC ensembles represent upcoming reward areas. (**A**) Bayesian decoding was used to identify representation of the reward areas as rats approached the choice point. Two examples are shown, from left and right trials, of decoded instances when ensemble firing patterns were more consistent with the upcoming reward area than the rat's actual position (gray circle). (**B**) The analyses of decoded spatial information focused on the two reward areas (black rectangles) but, importantly, was limited to time windows when the rat was located on the stem (red area). (**C**) Heat maps illustrating the average decoded probability from all of the 200 ms time bins used for decoding as the rat traversed the stem, with separate heat maps shown for each learning stage.

For illustration purposes, the left trials are mirror reversed so that all the data are illustrated with the correct goal location to the right. Note the faint clouds of probability at the reward areas (i.e. decoding to the reward areas) during the early learning stage. This becomes more prominent through late learning and only becomes selective for the correct reward area during overtraining. Stem locations are uniformly red because the decoding is most prevalent at the rat's actual current location on the stem. (**D**) Decoding to reward areas increased with training and surpassed chance levels (uniform non-local decoding, dashed line, see methods) only at criterion and during overtraining. Reward decoding was calculated by dividing the decoded probability that the rat was located in the rectangles of B by the expected probability as determined by relative area. Individual ensembles are plotted as open circles, while average reward area decoding for each training stage is plotted with error bars showing +/- SEM. (**E**) Selective decoding to the correct reward area only emerged (was statistically significant) during overtraining. Future reward decoding was defined as the normalized difference between decoding to the correct reward area and the opposite (incorrect) reward area ((p(correct) – p(incorrect)) / p(correct) + p(incorrect)). Individual ensembles are plotted as open circles, while average reward area decoding for each training stage is plotted with error bars showing +/- SEM. (**F**) Decoding to the correct reward area was not correlated with lateral position differences (summed across all four maze sectors, see **Figure 26A-B**) between stem traversals of left and right trials (top) or with running speed (bottom).

CHAPTER FOUR

Retrosplenial Ensembles Sustain Task-Related Activity Over a Delay

Abstract: Spatial working memory is an essential feature of goal directed navigation. However, it is not known whether the RSC, which plays an important role in spatial LTM, might also support this memory process. We trained rats on a continuous spatial alternation task and then introduced a delay between trials, thereby requiring rats to sustain a representation of the current trial type (go-left or go-right). Lesions of the RSC greatly impaired alternation performance on the delay version of the task. During the delay, individual RSC neurons maintained activity that was specific to the current trial type, and we were able to decode the rat's future choice (left or right turn) from the activity of RSC ensembles. Additionally, this trial-type specific activity was greatly reduced on error trials, when the rat failed to remember the correct response. This is the first direct evidence that the RSC contributes to spatial working memory processes.

Introduction

LTM supports many forms of cognition. The RSC, in particular, plays an important role in the long-term storage of spatial memories (Czajkowski et al., 2014; Katche et al., 2013), and is active during a range of processes that depend on the reactivation of LTM, such as autobiographical memory and future planning (Schacter & Addis, 2007; Spreng, Mar, & Kim, 2009) and the simulation of navigation routes (Brown et al., 2016). However, little research has considered how the RSC might contribute to spatial working memory, a related process whereby spatial information is held in memory over seconds (Spellman et al., 2015).

The vast majority of research on working memory has focused on the prefrontal cortex (PFC) in humans and monkeys because it shows strong neural activity during delay periods (Funahashi, Chafee, & Goldman-Rakic, 1993; Levy & Goldman-Rakic, 2000; Wilson, Scalaidhe, & Goldman-Rakic, 1993) and because individual PFC neurons sustain stimulus-specific activity over delays (Fuster & Alexander, 1971; Kubota & Niki, 1971; E. K. Miller, Erickson, & Desimone, 1996). Likewise, lesions of the PFC (Levy & Goldman-Rakic, 2000; Wilson et al., 1993) or other disruptions of PFC activity (Wegener, Johnston, & Everling, 2008) produce strong deficits in working memory tasks. However, recent work has additionally implicated posterior cortical regions, including the posterior parietal cortex (Crowe et al., 2013) and visual areas (E. K. Miller, Li, & Desimone, 1993; Motter, 1994), in the sustained representation of to-be-remembered items during the delay. In some cases, stimulus information can be more readily decoded from the activity of posterior regions than from PFC activity (Emrich, Riggall, Larocque, & Postle, 2013; Riggall & Postle, 2012). This has led some authors to suggest that the primary role of the prefrontal cortex in working memory is to reactivate stimulus representations stored in other cortical regions (Curtis & D'Esposito, 2003; Lara & Wallis, 2014; Postle, 2006).

The RSC is reciprocally interconnected with many regions known to play an important role in spatial working memory (Agster & Burwell, 2013; Burwell & Amaral, 1998; van Groen & Wyss, 1990, 1992, 2003) and makes critical contributions to spatial cognition through the encoding of contexts and spatial cues (for reviews see A. M. P. Miller, Vedder, Law, & Smith, 2014; Todd & Bucci, 2015). Some early evidence from lesion studies in rats has suggested that the RSC is also necessary to hold a spatial memory over a delay (Keene & Bucci, 2009; Nelson, Powell, Holmes, Vann, & Aggleton, 2015). However, nothing is known about how RSC ensembles might support this type of memory. To test this, we recorded activity in the RSC as rats performed a delayed spatial alternation task that required rats to hold a memory for 30 s at the beginning of each trial. We found that individual RSC neurons sustained activity during the delay that was specific to the rat's current behavioral goals. At the population level, we were able to decode the rat's future choice from the activity of RSC ensembles during the delay, and we found that this activity was greatly diminished during delay periods before the rat committed an error. This is the first direct evidence for an RSC role in working memory.

Results

RSC lesions impair delayed spatial alternation.

To test whether the RSC is required to hold a memory over a delay, we trained control rats and rats with neurotoxic (NMDA) lesions of the RSC on a modified t-maze designed for continuous spatial alternation (**Figure 28A**). The rats were rewarded with 0.2ml of chocolate milk for alternating between the left and right reward arms of the maze. Rats were given daily training sessions until they reached a criterion of 90% correct, followed by four additional asymptotic performance sessions. Rats were then trained on a delay version of the task, whereby the path to the reward was blocked at the beginning of each trial for 0 s (i.e. the block was removed as soon as the rat approached), 5 s, or 30 s. A mixed effects ANOVA comparing control and lesion group performance over the three delay lengths revealed a significant main effect of delay ($F(1,47) = 26.75$, $p < 0.001$) and a significant main effect of lesion ($F(1,47) = 13.05$, $p < 0.001$; **Figure 28B**), showing that rats with RSC lesions were impaired on the task. The delay by lesion interaction was marginally significant ($F(1,47) = 3.71$, $p = 0.06$), with a trend toward a larger lesion impairment at greater delay lengths. A planned comparison of the two lesion conditions at the 30 s delay length revealed an impairment among rats with RSC lesions at this delay ($t(15) = 2.93$, $p = 0.01$). The size of the lesions varied considerably (39.5 – 72.8% of the RSC, **Figure 28C**), and the impairment during the 30s delay was correlated with lesion size ($r = -0.78$, $p < 0.05$; **Figure 28C**).

RSC neurons distinguish between left and right trial types during the delay.

In order to examine the neurophysiological basis of RSC contributions to delay memory, we recorded from 123 neurons in 8 rats while they performed the 30 s delayed spatial alternation

task. Our recordings targeted Rgb bilaterally, although a small number of neurons from the RSA were also included. Previous studies have shown that RSC neurons encode spatial locations and direction of travel (Vedder et al, 2016; Alexander & Nitz, 2015), and that RSC neural firing can be used to determine the rat's future turn direction (Vedder et al, 2016). Consistent with this, we found that many RSC neurons showed spatially localized firing all along the path to the reward, and that this firing was specific to whether the rat was traveling to the left or right reward (**Figure 28D**). To quantify firing differences between left and right trials, we compared the firing of individual RSC neurons on each trial type (Left and Right) at four key periods during the trial: the reward (2 s after reward receipt), the approach (2 s before reward receipt), the stem (middle 1s during typical passes, see methods), and during the delay (30s before stem entrance; **Figure 28E,F**). Consistent with previous work on RSC spatial coding from our laboratory, we found that 64 RSC neurons (52.03%) distinguished between the left and right reward locations (as determined by a significant t-test of neural activity in the 1 s after reward receipt on correct left and right trials) and that 84 neurons (68.29%) differentiated between the left and right approach arms (2s before reward receipt). Importantly, we also found that 15 neurons (12.30%) differentiated between left and right trials while the rats traversed the stem of the maze (before turning left or right), far more than expected by chance ($p < 0.005$, as compared to a control distribution generated by shuffling left and right trial types 10,000 times). Lastly, we compared the activity of RSC single neurons during the 30 s delay period beginning when the rats entered the area at the base of the stem and ending when the block was removed and the rat entered the stem (**Figure 28A**). We found that 14 neurons (11.38%) differentiated between delay periods of left and right trials, which was more than expected by chance alone ($p < 0.005$, as compared to a distribution generated by shuffling left and right trial types). Examination of the firing rates of

these neurons revealed elevated firing for one context that was largely consistent over the course of the delay (**Figure 28G,H**), thereby indicating the current trial type throughout the delay period. This is the first neurophysiogical evidence for an RSC role retaining a memory over a delay.

RSC neuronal populations encode memory for trial type over a delay.

To test whether RSC neuronal populations encode the current trial type during delayed alternation, we compared population activity between left and right trials at each of the four key periods described above. To do this, we combined neurons from all rats into a single population, and then calculated mean firing rate vectors for each 250 ms time window during the session. We then visualized population activity over the entire session by plotting the first three principal components of this activity, and found that the resulting plot largely traced the outline of the maze (compare **Figure 29A** with **Figure 29B**), confirming that RSC activity is heavily modulated by spatial location. To additionally determine whether we could use RSC instantaneous activity to identify trial type, we employed a minimum distance classifier to assign activity from every 25 ms time window to one of the 8 trial-type specific periods (4 periods X 2 trial types). We found that we were able to accurately classify 69.29% of all time windows, far greater than expected by chance (binomial test, $p < 0.001$, **Figure 29C**). Classification was particularly accurate among the four periods, and most errors occurred between trial types within the periods that shared a spatial location (stem and delay). To therefore determine how well we could classify activity occurring during each period as belonging to either a left or right trial, we performed the classification procedure separately on each period, and found that all periods could be classified at a rate greater than chance (binomial tests, all $p < 0.001$; **Figure 29D**).

We next compared population activity during the delay periods of correct and error trials to determine whether superior delay activity preceded accurate alternation performance. To do this, we averaged population firing rate vectors across trials to form four average delay periods: correct left, correct right, error left, and error right. A plot of the first two principal components of correct trials illustrated that left and right neural activities remain distinct over the delay (**Figure 29E**). To then compare the quality of delay activity during correct and error trials, we computed a specificity index that measured how similar delay activity was to the mean of the correct trial type (e.g., left time windows to the left trial mean) compared to the incorrect trial type (e.g., left to right). Consistent with the idea that high quality delay activity precedes correct responses, we observed far higher specificity during the delay periods of correct trials than error trials ($t(115) = 11.57$, $p < 0.001$; **Figure 29F**). Interestingly, activity early in the delay period was highly specific during both correct and error trials (**Figure 29G**). However, while specificity remained similar throughout the delay during correct trials ($r = -0.06$, $p = 0.52$), specificity decreased dramatically during error trials ($r = -0.61$, $p < 0.001$).

Discussion

Despite a growing body of evidence showing that the RSC plays a crucial role in spatial cognition, next to nothing is known about how the RSC might support spatial working memory. We addressed this by testing the effects of RSC lesions on performance of a delayed alternation task, and by observing RSC ensemble activity while rats held a memory over the delay. Our findings point to an important role for the RSC in the sustained representation of information about the rat's upcoming behavioral response. In particular, we found that rats with RSC damage showed impaired delayed alternation performance, and that many neurons in the RSC showed sustained activity during the delay that was correlated with the rat's future choice. At the population level, we demonstrated that the failure of RSC ensembles to maintain the correct activity state was associated with a failure to select the correct reward at the end of the delay. These task-related signals during the delay occurred in addition to RSC ensemble coding for the rat's current spatial location and for time within the session. Together, these results suggest that the sustained activation of memory information stored in the RSC may be an important part of the brain's spatial working memory process.

Although the RSC is known to contribute to a range of episodic and working memory type processes in humans (Spreng et al., 2009), working memory research in animals has focused far more on the PFC due to its robust neural activity during delay periods (Fuster & Alexander, 1971; Kubota & Niki, 1971). Here we discovered that the RSC makes similar contributions to spatial working memory in rats, showing for the first time that the RSC sustains trial-specific activity over a delay, and that the disruption of this information is associated with memory failure. Our finding that 11% of RSC cells showed a preference for one trial type during the delay was similar to what has been found in the rat mPFC (thought to be comparable to the

dorsolateral PFC in monkeys; Seamans, Lapish, & Durstewitz, 2008), where 10% of neurons showed trial-type specificity on similar task (Yang, Shi, Wang, Peng, & Li, 2014). Nonetheless, important coding differences exist between the two regions. While our RSC neurons showed relatively persistent trial-type specific activity over the course of the delay, mPFC neurons showed trial-type specificity in punctuate bursts during the period, perhaps as part of a "relay-race" coding strategy (Yang et al., 2014). Additionally, while mPFC ensembles appear to encode the future reward location during the delay (Spellman et al., 2015), RSC ensembles appeared to encode only general trial-type information.

Many RSC neurons also fired on the stem of the maze according to whether the rat would soon visit the left or right reward. However, the proportion of neurons showing this effect was far lower than we previously observed on a continuous spatial alternation task (**Figure 25B**). Interestingly, a similar decrease was also observed in the hippocampus (i.e. fewer trial-type specific neurons in the delay task; Ainge, van der Meer, Langston, & Wood, 2007; Wood, Dudchenko, Robitsek, & Eichenbaum, 2000). This may be due to the delay period disrupting the rat's smooth trajectory from one reward to the other, and thereby reducing the influence of recent spatial positions and motor behaviors on stem firing (Alexander & Nitz, 2015; Cho & Sharp, 2001). Alternatively, this effect might simply reflect the animal's weaker memory for the current trial type during the more-difficult delay task. We have previously found that trial-type specificity was decreased on days when the rat performed less accurately (**Figure 25F,G**). Furthermore, we saw a similar proportions of neurons with trial-type specific firing during the delay task as we did during learning on the continuous task, when performance was similar.

A role for the RSC in working memory is consistent with the greater role of the neocortex supporting cognition by reactivating LTM representations. Much work has shown that many

memories that are initially dependent on the hippocampus are later consolidated into neocortical regions such as the RSC for long-term storage (Katche et al., 2010; Katche et al., 2013), and that the hippocampus supports recollection by reactivating this LTM (Takata et al., 2015). This reactivating function may explain why the hippocampus is specifically involved in stimulus encoding during working memory tasks (i.e. it may be reactivating an existing long-term representation of the stimulus; Spellman et al., 2015), while the PFC (along with posterior cortical regions Pasternak & Greenlee, 2005) is required to sustain the representation. Our work here suggests that the RSC may be importantly involved in this process, either as a target of hippocampal reactivation and PFC maintenance (or a combination of the two; Ito, Zhang, Witter, Moser, & Moser, 2015), or by supporting memory maintenance more directly, such as by managing attention to irrelevant (Ng, Noblejas, Rodefer, Smith, & Poremba, 2007) or otherwise surprising information (Pearson, Heilbronner, Barack, Hayden, & Platt, 2011).

Methods

Subjects

Subjects were 28 adult male Long Evans rats (Charles River Laboratories, Wilmington, MA) weighing 250g-300g upon arrival. Eight rats were used in the electrophysiology study and 20 rats were used in the lesion study. Of the 10 rats that received RSC lesions, 2 were excluded from the analysis due to hippocampal damage, and 1 was excluded because the RSC damage was unilateral. Rats were placed on a 12hr/12hr light/dark cycle with lights on at 8am and allowed to acclimate to the vivarium for at least one week prior to surgery. After surgery, rats were placed on food restriction until they reached 80-85% of their free-feeding weight. Water was always available ad libitum. The rats in both studies were previously used in a separate learning study and the data from that study is reported in chapter 2. All procedures complied with the guidelines of the Cornell University Animal Care and Use Committee.

Surgery

Lesions. Twenty rats were anethetized with isoflorane gas (1-5% in oxygen) and placed in a Kopf stereotaxic apparatus. The skin was retracted and holes were drilled through the skull above each of the injection sites. Ten rats received bilateral neurotoxic (N-methyl-D-aspartate [NMDA], 10µg/ml) lesions of the RSC. NMDA was injected by hand in volumes of 0.20 – 0.35 µl using a custom-made glass injection canula (100µm diameter) attached to a Hamilton Syringe by sterile plastic tubing. The stereotaxic coordinates and injection volumes were:

1. 0.35µL at -2.2 (AP), ±0.5 (ML), -3.0 (DV)
2. 0.35µL at -3.9 (AP), ±0.5 (ML), -3.0 (DV)

3. 0.20µL at -5.5 (AP), ±0.5 (ML), -3.5 (DV)

4. 0.35µL at -5.5 (AP), ±1.0 (ML), -2.8 (DV)

5. 0.35µL at -6.7 (AP), ±1.1 (ML), -2.8 (DV)

6. 0.30µL at -8.0 (AP), ±1.3 (ML), -2.8 (DV)

Coordinates were taken from Bregma (AP), the midline (ML), and the surface of the skull (DV), respectively. The injection cannula was left in place 1 min before and 5 min after each infusion. An additional ten rats received sham lesions of the RSC consisting of lowering the injection cannula into the brain but not injecting NMDA.

Electrophysiology. Twelve rats had a custom-built electrode microdrive implanted, which contained 20 moveable tetrodes (16 recording tetrodes and 4 reference tetrodes) made from twisting four 17µm platinum/iridium (90%/10%) wires, platinum plated to an impedance of 100-500kΩ, and arranged in two 10-tetrode linear arrays (one in each hemisphere) that spanned approximately 5mm along the rostrocaudal axis of the brain. Tetrodes were stereotaxically positioned bilaterally just beneath the cortical surface (2-7mm posterior to Bregma, ±1.5mm lateral) with the tetrodes angled 30 degrees toward the midline. Rats were given 7 days to recover from surgery prior to lowering the tetrodes into the RSC (35-70µm daily) over the course of several days until a depth of at least 1mm was reached in order to ensure that the tetrodes were in the granular b subregion (discussed below).

Apparatus

The behavioral apparatus was a black PVC continuous T-maze (120 cm long stem x 100 cm wide x 68 cm above the floor) with metal reward cups embedded in the ends of the arms. Chocolate milk (0.2 ml, Nestle's Nesquik) could be delivered to the reward cups via an elevated

reservoir controlled by solenoid valves activated by foot-pedal switches. The maze was located in the center of a circular arena enclosed by black curtains with visual cues of various shapes, sizes, and colors. The room was illuminated by a ring of LED lights around the edge of the ceiling. A continuous background masking noise was played from a speaker located directly above the maze.

Behavioral Training Procedures

Prior to training, rats were acclimated to the maze and chocolate milk rewards with daily periods of free exploration on the maze until rats consumed 20 rewards within the first 10min of an acclimation session (mean = 4.5 acclimation days). After acclimating to the maze, rats were trained on a continuous spatial alternation task in which the rats were rewarded only if they approached the opposite (left or right) reward location from the previous trial. Both cups were baited on the first trial. Entries into the same arm as the previous trial were scored as an error and were not rewarded. Rats were gently ushered back if they left the continuous alternation route. Rats were not allowed to correct their errors. Rats were given 40 trials/day until they achieved a criterion of 90% correct on two consecutive days. After achieving this criterion, rats were then trained on a delay version of the task whereby the path to the reward was blocked (black wooden block, 18” X 8” X 2”, attached to a plastic base that fit into the stem of the maze) at the base of the stem at the beginning of each trial. In the lesion experiment, delay lengths were either 0 s (block was removed as soon as the rat approached), 5 s, or 30 s. In the electrophysiology experiment, every trial during the delay sessions had a 30 s delay. The experimenter monitored delay lengths with a stopwatch.

Recordings

Neuronal spike data and video data were collected throughout learning (Digital Cheetah Data Acquisition System, Neuralynx, Inc. Bozeman, MT), filtered at 600Hz and 6kHz, digitized and stored to disc along with timestamps for offline sorting (SpikeSort3D, Nueralynx, Inc.). The rat's position was monitored by digitized video of an LED array attached to the rat's head. The time of reward receipt was measured with a grounding circuit that detected oral contact with the chocolate milk reward.

Histology

After completion of the experiment, rats were transcardially perfused with 4% paraformaldehyde in phosphate buffered saline. Brains were removed and stored for at least 24hrs in 4% paraformaldehyde before being transferred to a 30% sucrose solution for storage until slicing. Coronal sections (40µm) were stained with 0.5% cresyl violet for visualization of tissue damage (in the case of NMDA lesions) or tetrode tracks (for electrophysiology recording implants). Tissue damage was quantified by laying a grid (250µm to-scale grid spacing) over an enlarged image of the stained tissue and dividing the number of grid intersections located over damaged RSC areas by the number of intersections located over the entire RSC. Boundaries of the RSC were determined in accordance with The Rat Brain in Stereotaxic Coordinates(Paxinos & Watson, 1998). Tetrode positions were identified using depth records noted during tetrode lowering and tracks observed in the stained tissue (**Figure 21**). Neuronal records from tetrodes located outside of the RSC were excluded from the data set. As in our previous work(Vedder, Miller, Harrison, & Smith, 2016), our recordings targeted the granular b subregion of the RSC, although small numbers of neurons from the granular a subregion or the dysgranular RSC were

also included. There were no conspicuous differences in the firing properties of neurons recorded in different subregions, different hemispheres, or at different AP coordinates.

Data analysis

Consistent with previous work from our lab, we found that RSC ensembles drifted steadily over the course of the recording session (**Figure 30**; compare with **Figure 14A**). To better isolate the effects of interest, we statistically removed the effect of population drift by fitting a line to the relationship between time and firing rate independently for each cell, and then subtracting the value of the fit line from the firing rates at each time point. However, similar results were obtained for all analyses when this control was not performed.

We examined RSC neuronal activity during four task periods: (1) the delay, (2) as the rat traversed the stem of the maze (before turning left or right), (3) the approach to the reward (as the rat ran along the arms of the maze), and (4) the receipt of the reward (defined as the 2 s after the first lick detection). For each of these periods, individual neurons were classified as having a trial-type specific response by comparing firing rates on correct left trials with firing rates on correct right trials using a t-test. We then determined whether the proportion of neurons showing trial-type specific firing was greater than expected by chance alone by randomly relabeling correct trials as either left or right individually for each neuron 10,000 times, and after each time recalculating the proportion of neurons meeting the classification criterion. The observed proportion was considered statistically significant if it was greater than 97.5% of the shuffled proportions.

To determine whether RSC ensemble firing could be used to distinguish between the task periods of each trial type (4 periods X 2 trial types), we employed a minimum distance classifier

to sort instantaneous RSC activity into one of the eight trial-type periods. To do this, we combined neurons across rats and sessions to form ensemble firing rate vectors for every 250 ms time window during visits to each of the eight periods. We then used an iterative procedure whereby we excluded one time window from the data set, calculated mean firing rate vectors for every period from the remaining time windows, and then computed the standardized Euclidean distance between the excluded time window and the eight means. The activity from the excluded time window was then assigned to a period according to which it was closest to in neural state space (i.e. the minimum distance). In order to further determine whether RSC ensemble firing could be used to distinguish between trial types during each period, we repeated the above procedure independently for each period, classifying instantaneous firing into either left or right trial types. In all cases, we determined whether the proportion of successful classifications was greater than expected by chance alone by randomly relabeling time window classification categories (either 8 periods or 2 trial types) 10,000 times, and after each time recalculating the proportion of successful classifications. The observed proportion was considered statistically significant if it was greater than 97.5% of the shuffled proportions.

To compare RSC activity during delay periods on correct and error trials, we first identified all sessions where the rat committed at least five left errors (i.e. the rat went to the left when he should have gone to the right) and all sessions where the rat committed at least five right errors. We then used an iterative procedure whereby we excluded one trial from the data set, calculated mean left and right firing rate vectors the from remaining correct trials, and then computed the standardized Euclidean distance from each time window in the excluded trial to the two correct-trial means. To ensure that differences between correct and error trials could not be due to greater sampling of correct trials (which occurred more often), we included only five

correct left trials and five correct right trials drawn from the middle of the session. A trial-type specificity measure was then calculated from the distances such that, for a left trial time window, specificity was equal to the distance to the mean of left trials minus the distance to the mean of right trials, divided by the sum of the distances. All specificity calculations were performed separately on the sessions with left errors and sessions with right errors, and then the outcomes were averaged across them. Specificity calculations were then averaged across left and right correct trials and across left and right error trials.

Acknowledgements

I am particularly grateful for the contributions of Hei Jun Li, who completed his undergraduate thesis while assisting on this project. I also thank William Mau, Sarah Parauda, Keunhyung Yu, and Alexandra Tse for assistance with animal training and electrode hyperdrive fabrication.

References

Agster, K. L., & Burwell, R. D. (2013). Hippocampal and subicular efferents and afferents of the perirhinal, postrhinal, and entorhinal cortices of the rat. *Behav Brain Res, 254*, 50-64. doi: 10.1016/j.bbr.2013.07.005

Ainge, J. A., van der Meer, M. A., Langston, R. F., & Wood, E. R. (2007). Exploring the role of context-dependent hippocampal activity in spatial alternation behavior. *Hippocampus, 17*(10), 988-1002. doi: 10.1002/hipo.20301

Alexander, A. S., & Nitz, D. A. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nat Neurosci, 18*(8), 1143-1151. doi: 10.1038/nn.4058

Brown, T. I., Carr, V. A., LaRocque, K. F., Favila, S. E., Gordon, A. M., Bowles, B., . . . Wagner, A. D. (2016). Prospective representation of navigational goals in the human hippocampus. *Science, 352*(6291), 1323-1326. doi: 10.1126/science.aaf0784

Burwell, R. D., & Amaral, D. G. (1998). Cortical afferents of the perirhinal, postrhinal, and entorhinal cortices of the rat. *J Comp Neurol, 398*(2), 179-205.

Cho, J., & Sharp, P. E. (2001). Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behav Neurosci, 115*(1), 3-25.

Crowe, D. A., Goodwin, S. J., Blackman, R. K., Sakellaridi, S., Sponheim, S. R., MacDonald, A. W., 3rd, & Chafee, M. V. (2013). Prefrontal neurons transmit signals to parietal neurons that reflect executive control of cognition. *Nat Neurosci, 16*(10), 1484-1491. doi: 10.1038/nn.3509

Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends Cogn Sci, 7*(9), 415-423.

Czajkowski, R., Jayaprakash, B., Wiltgen, B., Rogerson, T., Guzman-Karlsson, M. C., Barth, A. L., . . . Silva, A. J. (2014). Encoding and storage of spatial information in the retrosplenial cortex. *Proc Natl Acad Sci U S A, 111*(23), 8661-8666. doi: 10.1073/pnas.1313222111

Emrich, S. M., Riggall, A. C., Larocque, J. J., & Postle, B. R. (2013). Distributed patterns of activity in sensory cortex reflect the precision of multiple items maintained in visual short-term memory. *J Neurosci, 33*(15), 6516-6523. doi: 10.1523/jneurosci.5732-12.2013

Funahashi, S., Chafee, M. V., & Goldman-Rakic, P. S. (1993). Prefrontal neuronal activity in rhesus monkeys performing a delayed anti-saccade task. *Nature, 365*(6448), 753-756. doi: 10.1038/365753a0

Fuster, J. M., & Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science, 173*(3997), 652-654.

Ito, H. T., Zhang, S. J., Witter, M. P., Moser, E. I., & Moser, M. B. (2015). A prefrontal-thalamo-hippocampal circuit for goal-directed spatial navigation. *Nature, 522*(7554), 50-55. doi: 10.1038/nature14396

Katche, C., Bekinschtein, P., Slipczuk, L., Goldin, A., Izquierdo, I. A., Cammarota, M., & Medina, J. H. (2010). Delayed wave of c-Fos expression in the dorsal hippocampus involved specifically in persistence of long-term memory storage. *Proc Natl Acad Sci U S A, 107*(1), 349-354. doi: 10.1073/pnas.0912931107

Katche, C., Dorman, G., Gonzalez, C., Kramar, C. P., Slipczuk, L., Rossato, J. I., . . . Medina, J. H. (2013). On the role of retrosplenial cortex in long-lasting memory storage. *Hippocampus, 23*(4), 295-302. doi: 10.1002/hipo.22092

Keene, C. S., & Bucci, D. J. (2009). Damage to the retrosplenial cortex produces specific impairments in spatial working memory. *Neurobiol Learn Mem, 91*(4), 408-414. doi: 10.1016/j.nlm.2008.10.009

Kubota, K., & Niki, H. (1971). Prefrontal cortical unit activity and delayed alternation performance in monkeys. *J Neurophysiol, 34*(3), 337-347.

Lara, A. H., & Wallis, J. D. (2014). Executive control processes underlying multi-item working memory. *Nat Neurosci, 17*(6), 876-883. doi: 10.1038/nn.3702

Levy, R., & Goldman-Rakic, P. S. (2000). Segregation of working memory functions within the dorsolateral prefrontal cortex. *Exp Brain Res, 133*(1), 23-32. doi: 10.1007/s002210000397

Miller, A. M. P., Vedder, L. C., Law, L. M., & Smith, D. M. (2014). Cues, context, and long-term memory: the role of the retrosplenial cortex in spatial cognition. *Frontiers in Human Neuroscience, 8*. doi: 10.3389/fnhum.2014.00586

Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci, 16*(16), 5154-5167.

Miller, E. K., Li, L., & Desimone, R. (1993). Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci, 13*(4), 1460-1478.

Motter, B. C. (1994). Neural correlates of feature selective memory and pop-out in extrastriate area V4. *J Neurosci, 14*(4), 2190-2199.

Nelson, A. J., Powell, A. L., Holmes, J. D., Vann, S. D., & Aggleton, J. P. (2015). What does spatial alternation tell us about retrosplenial cortex function? *Front Behav Neurosci, 9*, 126. doi: 10.3389/fnbeh.2015.00126

Ng, C. W., Noblejas, M. I., Rodefer, J. S., Smith, C. B., & Poremba, A. (2007). Double dissociation of attentional resources: prefrontal versus cingulate cortices. *J Neurosci, 27*(45), 12123-12131. doi: 10.1523/jneurosci.2745-07.2007

Pasternak, T., & Greenlee, M. W. (2005). Working memory in primate sensory systems. *Nat Rev Neurosci, 6*(2), 97-107. doi: 10.1038/nrn1603

Paxinos, G., & Watson, C. (1998). *The rat brain in stereotaxic coordinates*.

Pearson, J. M., Heilbronner, S. R., Barack, D. L., Hayden, B. Y., & Platt, M. L. (2011). Posterior cingulate cortex: adapting behavior to a changing world. *Trends Cogn Sci, 15*(4), 143-151. doi: 10.1016/j.tics.2011.02.002

Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience, 139*(1), 23-38. doi: 10.1016/j.neuroscience.2005.06.005

Riggall, A. C., & Postle, B. R. (2012). The relationship between working memory storage and elevated activity as measured with functional magnetic resonance imaging. *J Neurosci, 32*(38), 12990-12998. doi: 10.1523/jneurosci.1892-12.2012

Schacter, D. L., & Addis, D. R. (2007). Constructive memory: The ghosts of past and future. *Nature, 445*(7123), 27-27.

Seamans, J. K., Lapish, C. C., & Durstewitz, D. (2008). Comparing the prefrontal cortex of rats and primates: insights from electrophysiology. *Neurotox Res, 14*(2-3), 249-262. doi: 10.1007/bf03033814

Spellman, T., Rigotti, M., Ahmari, S. E., Fusi, S., Gogos, J. A., & Gordon, J. A. (2015). Hippocampal-prefrontal input supports spatial encoding in working memory. *Nature, 522*(7556), 309-314. doi: 10.1038/nature14445

Spreng, R. N., Mar, R. A., & Kim, A. S. (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*(3), 489-510. doi: 10.1162/jocn.2008.21029

Takata, N., Yoshida, K., Komaki, Y., Xu, M., Sakai, Y., Hikishima, K., . . . Tanaka, K. F. (2015). Optogenetic activation of CA1 pyramidal neurons at the dorsal and ventral hippocampus evokes distinct brain-wide responses revealed by mouse fMRI. *PLoS One, 10*(3), e0121417. doi: 10.1371/journal.pone.0121417

Todd, T. P., & Bucci, D. J. (2015). Retrosplenial Cortex and Long-Term Memory: Molecules to Behavior. *Neural Plast, 2015*, 414173. doi: 10.1155/2015/414173

van Groen, T., & Wyss, J. M. (1990). Connections of the retrosplenial granular a cortex in the rat. *J Comp Neurol, 300*(4), 593-606. doi: 10.1002/cne.903000412

van Groen, T., & Wyss, J. M. (1992). Connections of the retrosplenial dysgranular cortex in the rat. *J Comp Neurol, 315*(2), 200-216. doi: 10.1002/cne.903150207

van Groen, T., & Wyss, J. M. (2003). Connections of the retrosplenial granular b cortex in the rat. *J Comp Neurol, 463*(3), 249-263. doi: 10.1002/cne.10757

Vedder, L. C., Miller, A. M., Harrison, M. B., & Smith, D. M. (2016). Retrosplenial Cortical Neurons Encode Navigational Cues, Trajectories and Reward Locations During Goal Directed Navigation. *Cereb Cortex*. doi: 10.1093/cercor/bhw192

Wegener, S. P., Johnston, K., & Everling, S. (2008). Microstimulation of monkey dorsolateral prefrontal cortex impairs antisaccade performance. *Exp Brain Res, 190*(4), 463-473. doi: 10.1007/s00221-008-1488-4

Wilson, F. A., Scalaidhe, S. P., & Goldman-Rakic, P. S. (1993). Dissociation of object and spatial processing domains in primate prefrontal cortex. *Science, 260*(5116), 1955-1958.

Wood, E. R., Dudchenko, P. A., Robitsek, R. J., & Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron, 27*(3), 623-633.

Yang, S. T., Shi, Y., Wang, Q., Peng, J. Y., & Li, B. M. (2014). Neuronal representation of working memory in the medial prefrontal cortex of rats. *Mol Brain, 7*, 61. doi: 10.1186/s13041-014-0061-2

Figures

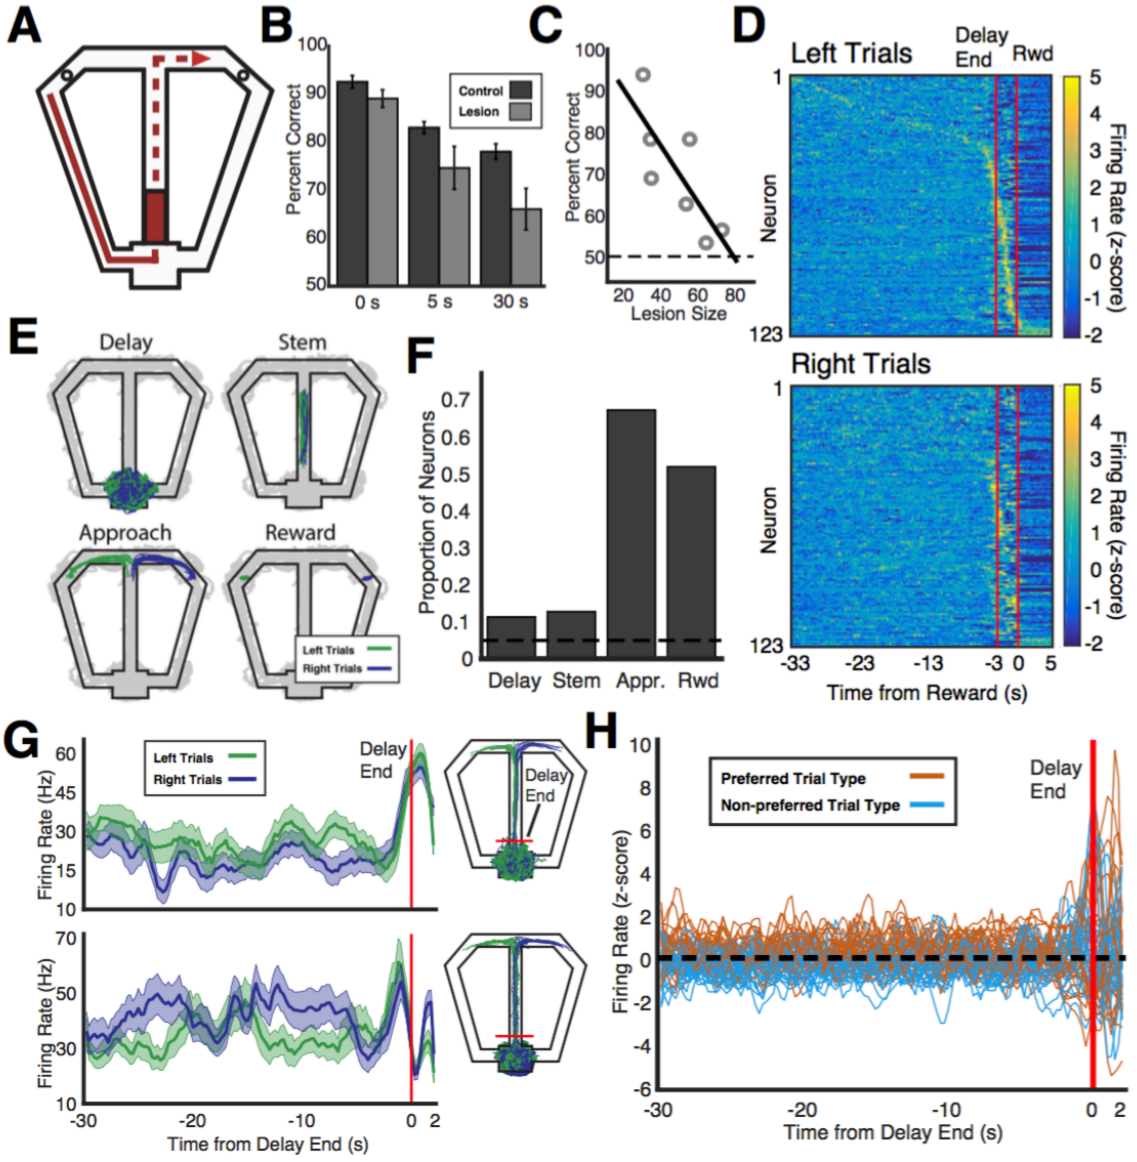

Figure 28 Single neurons in the RSC encode space and trial type. (**A**) Schematic diagram of the delayed t-maze alternation task. After visiting one of the two reward locations (circles), rats returned to the stem and had to approach the opposite location for reward. A block (red square) prevented the rat from entering the stem for 30 s at the beginning of each trial. (**B**) Permanent lesions of the retrosplenial cortex impair delayed spatial alternation performance. Behavioral performance is plotted for the three delay durations. (**C**) The correlation between performance

and lesion size is shown for lesion rats on the 30 s delay task. Chance performance is indicated by the dotted line. (**D**) All recordings employed the 30 s delay task. Standardized firing rates are plotted for left and right trials for each neuron, sorted by when the firing peak occurred on left trials. Most peaks occurred at the end of the delay (approximate delay end, first red line) or while the rat ran toward the reward (instant of first lick, second red line). Some cells showed similar peak times on left and right trials, but most did not, indicating that task-related firing differed between trial types. (**E**) We compared firing rates at four periods during each trial. The spatial positions of one rat during one session are shown for each of the four periods. Note that the rat is occupying the same spatial positions on left and right trials during the delay and stem periods (i.e. the blue and green paths overlap), but is occupying different positions during the approach and reward periods. Grey shading indicates all visited locations. (**F**) The proportion of neurons whose firing rates on left trials were different from right trials is shown for each of the four trial periods. All proportions were significantly greater from chance (dotted line, 0.05). (**G**) Two example neurons with distinct firing on left and right trials during the delay period are shown (left). Lines show mean firing rate over all left and right trials +/- SEM. Corresponding spatial positions during were highly overlapping during the delay (right). (**H**) Trial type was discernable from the firing of many neurons. Mean firing rates on are shown for all neurons with significantly different firing during left and right delay periods, colored by the higher firing (i.e. preferred) trial type.

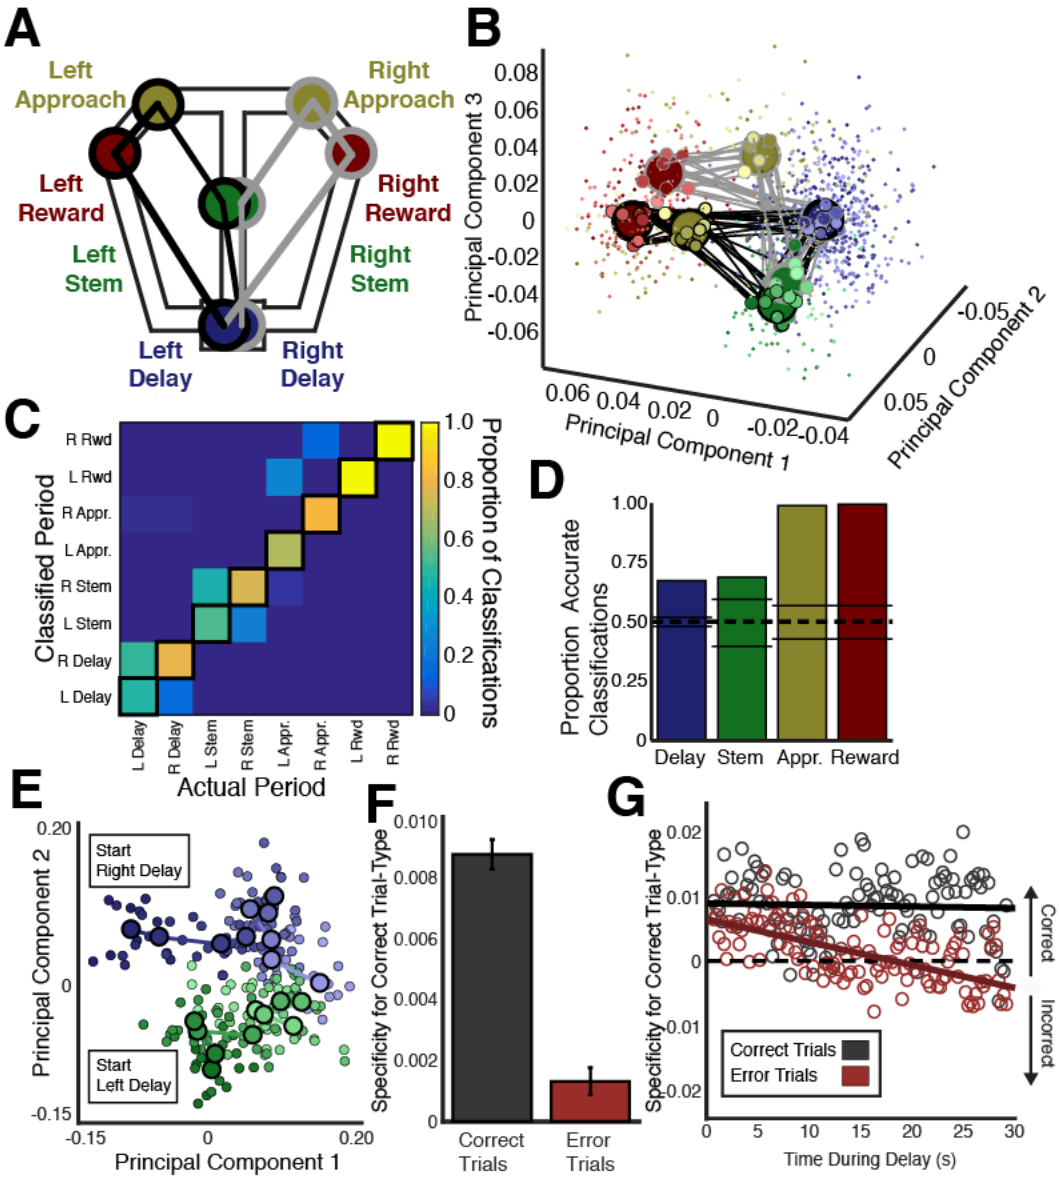

Figure 29 RSC ensembles encode space and trial type. (**A**) Diagram showing the maze locations corresponding to the eight trial type specific periods (four trial periods X two trial types). (**B**) Three dimensional plot of the first three principal components of RSC ensemble activity (combined across subjects and sessions, see methods) collapsed across time during session. For a non-collapsed version, see Figure 3. Colors are as in A. Small dots show activity from individual 250 ms time windows, Medium dots show average activity from individual trials, large dots show average activity during that trial-type specific period. Note that RSC ensemble activity in

state-space largely recapitulates the spatial relationships shown in A. (**C**) Classification accuracy of all time window activity into the eight periods is shown. Each column shows the proportion of time windows taken during that "actual" period that were classified into each of the eight "classified" periods (rows). The highest proportion was always to the correct period (black outlines). Classification errors were most common between periods that shared spatial locations. (**D**) Classification accuracy of all time window activity into left and right trial types is shown. Each bar shows the proportion of time windows taken during that period that were classified into the correct trial type. All accuracies were greater than chance (dotted line; chance range indicated by solid lines). (**E**) The first two principal components from average correct left (green) and right (blue) trials are plotted. Small dots show individual 250 ms time windows, while large dots show mean activity over a 3 s window. Lighter shading indicates later times during the delay. Note how delay activity from correct trials remains segregated by trial type over the delay. (**F**) RSC neural ensemble activity was less trial-type specific during error trials. Ensemble activity from individual delay periods is plotted in terms of trial-type specificity, defined as the difference between the distances to the representations of the correct (e.g., if the previous trial was left, then the correct trial is right) and incorrect trial types, normalized by the sum of both distances. Bars show means +/- SEM. (**G**) Colored dots show the specificity of ensemble activity from correct and error trials (as in F) plotted in terms of time during the average delay period. Points along the dotted line are equidistant to both correct and incorrect trial types, while points further from the dotted line indicate stronger ensemble preferences for one trial type over the other. Note that ensemble activity from correct trials reliably prefers the correct trial type over the delay. However, ensemble activity from error trials shows decreased preference for the correct trial type toward the end of the delay.

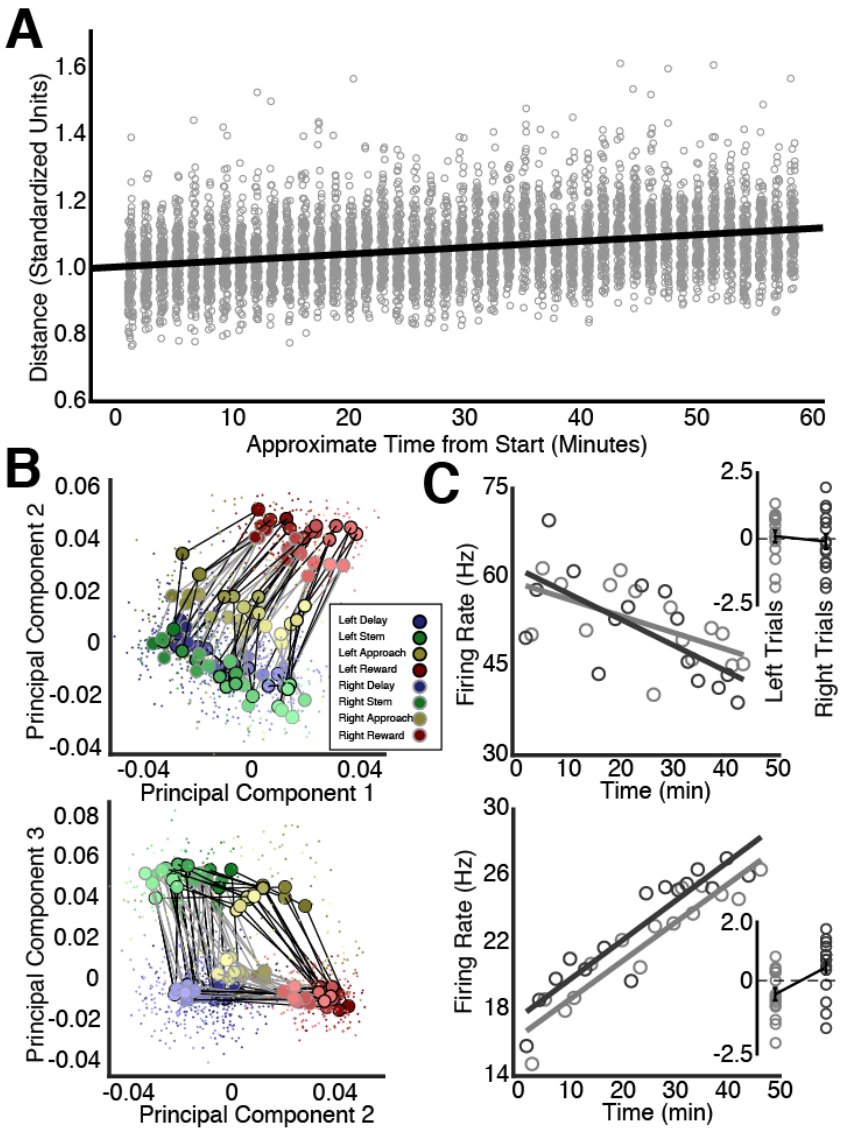

Figure 30 RSC ensembles show population drift over time. (**A**) Grey circles show RSC ensemble activity (combined across subjects and sessions, see methods) from each time window (during the four trial periods) plotted in terms of its standardized distance from the mean of ensemble activity during the first minute of the session. Note how the RSC ensemble moves steadily away from its starting state over the course of the session. The correlation between distance and time was highly significant (black line, $r = 0.31$, $p < 0.001$). (**B**) Plots of the first three principal components of RSC ensemble activity show the stability of trial period encoding

(colored dots) despite the steady population drift. (**C**) Two example neurons that show significant relationships between firing rate during the delay and time during the overall session (i.e. drift). Some neurons showed drift without encoding trial-type (Top). However, other neurons could be shown to encode trial type after controlling for drift (bottom inset, see methods). Circles show rates during the delay on a single trial, colored by trial type. Lines model drift within that trial type. Insets show the z-scored rates for each trial type after controlling for drift.

CHAPTER FIVE

The Predictive Engram

Context is a critical mechanism for guiding and organizing behavior. Recent research on neocortical LTM suggests that the same associative memories that support context representations also provide information about future events. This is similar to the tradition of understanding sensory perception as a generative process whereby sensory representations are produced internally based on predictions about future input. In this view, the brain models the underlying reality of its environment to best explain its limited experience, and then evokes that model to make predictions about future experiences. Here, I argue that this process is made possible in the memory system via interactions between the hippocampus and the neocortex. In particular, I suggest that most context memories are a product of an iterative learning process whereby (1) LTM context associations become active when an animal re-enters the context, (2) the activation of these memories gives the animal useful information about the environment, (3) this information improves learning and other processes in that environment, (4) newly learned information is automatically accommodated (i.e. reconsolidated) into the active context memory, (5) this new learning improves the context memory, thereby reinforcing the entire process.

The Exploitation of Redundancy

Behaving organisms are trapped between two undesirable truths: at no instant do they possess all of the information about their environment, but, at every instant, they must interact with that environment. To bridge the divide, organisms have developed strategies for maximizing the amount of information they can take in, and for extrapolating beyond this limited information to maximize their efficiency. These strategies rely on exploiting the tendency of

environments to repeat—to be redundant—over both space and time. For example, in order to maximize the amount of information that can be encoded from natural scenes, the visual system employs a coding strategy based upon the principle of redundancy reduction (Horace Barlow, 1961; H. Barlow, 2001), whereby redundancy among the many sensory receptors tasked with encoding the incoming sensory signal is recoded in order to encode the most information in the fewest number of cells. This is made possible because natural scenes contain statistical regularities that clearly distinguish them from randomness (Field, 1987). Broadly stated, redundancy reduction is a method for exploiting these regularities by assuming information that can be determined independently, thereby reducing the processing required and increasing the total amount of information that can be encoded.

Similar principles are employed in higher processing regions such as the neocortex, where information about previous sensory experiences is stored. By considering information about stimuli and events that have been encountered before, cortical regions can extrapolate beyond the information of the instant to a more general model of the sensory world (H. Barlow, 2001). This changes the question from "how much can I encode about this image?" to "what do I know about this images like this one?" while maintaining a processing approach that is similar to the redundancy reduction employed by earlier sensory regions. One characterization of this processing scheme highlights the importance of generative, top-down memory signals, such as the use of contextual information to estimate the probability of underlying identities (i.e. "Given everything I can see, what might this be an image of?"; (Fenske, Aminoff, Gronau, & Bar, 2006; Yuille & Kersten, 2006).

This framework emphasizes the importance of long-term memory for identifying the associations between different stimuli that can then be used to judge the likelihood of future

events. These associations, once reactivated, constitute a prediction about the likelihood of other stimuli and events. As with redundancy reduction, these predictions are made possible because the events and stimuli that comprise environments are not distributed randomly throughout space and time, but instead naturally group in distinct environments and/or circumstances. The activation of the appropriate associations is a kind of redundancy reduction that capitalizes on the tendency of nature to produce similar outcomes under similar conditions. Although the redundancy here is across time and space (as opposed to a 2D scene), the outcome is the same: the organism knows more despite encoding less.

Context Representations

Here, I use context to mean the brain-wide instantaneous predictions about the likelihood of stimuli and events. This intentionally includes various context-like phenomena such as schemas, semantic networks, priming, and other phenomena that highlight the importance of memory for inferring relationships or making predictions. The underlying idea behind this definition is that stimuli are non-randomly distributed in space and time, and that, in detecting and learning about the co-occurrence of stimuli, we form long-term associations that enable contextual predictions. Once formed, these associations naturally lend themselves to the prediction of future events through various cognitive processes (e.g., spread of activation, pattern completion) whereby an input stimulus leads to the reactivation of an associative network.

The reactivation of the associative network explains many of the phenomena described as contextual memory, such as the improvement in memory expression observed when retrieval conditions are similar to those present at encoding (Godden & Baddeley, 1975; S. M. Smith & Vela, 2001). A classic example of this phenomenon occurs in subjects who learn a list of stimuli

in a room with black walls will later recall more of the training stimuli if they are tested in the same black room as opposed to a white room (e.g., S. Smith, Glenberg, & Bjork, 1978). In terms of the process described above, the list of stimuli are hypothesized to develop associations with the many sensory features of the room (e.g., visual input from the black walls) during training, such that any future activation of those sensory features (e.g., seeing the black walls again) will also reactivate the list of stimuli in memory, thereby improving their retrieval.

The retrieval of contextually related information can also affect encoding and learning. In ambiguous instances, such as when the stimulus is too complex to fully encode, subjects often falsely remember stimuli that are consistent with the presentation context. Perhaps the most basic demonstration of this can be seen when subjects attempt to memorize a list of words with a semantic theme. For example, when presented with a series of words such as "tired," "nap," "bed," etc., subjects are more likely to later falsely report having seen a related word, such as "sleep," than an unrelated word, such as "car," even though neither of these words are presented, and the subject is explicitly warned against guessing (Deese, 1959; Roediger & McDermott, 1995). Similarly, after a brief exposure to an empty office, many subjects incorrectly report having seen objects such as books, a filing cabinet, and a coffee cup—all items associated with an office but not actually present (Brewer & Treyens, 1981).

These contextually related additions to memory are often described as false memories. However, the memory system can—and usually does—benefit from this type of context information, such as when memory for a passage is improved by an informative title (Bransford & Johnson, 1972), or when memory for a spatial arrangement is improved by the subject's prior experience with similar layouts (Chase & Simon, 1973; Tse et al., 2007). It makes sense that these associations are helpful than hurtful, as the most frequent co-occurrences in the

environment provide the greatest opportunity to strengthen their associations. Therefore, when these associations are reactivated, they naturally predict the most frequent (i.e. probable) outcomes.

Hippocampal-neocortical Interactions Supporting Context Processing

The primary candidate for representing context in the brain is the hippocampus. The hippocampus contributes importantly to associative processing (Komorowski, Manns, & Eichenbaum, 2009; Talk, Gandhi, & Matzel, 2002), spatial representations (Morris, Garrud, Rawlins, & O'Keefe, 1982; O'Keefe & Dostrovsky, 1971), and many forms of contextual memory (Anagnostaras, Gale, & Fanselow, 2001; Holland & Bouton, 1999). Moreover, individual neurons in the hippocampus suddenly and collectively alter their activity when rats move between putative contexts, such as distinct spatial environments (Muller & Kubie, 1987), internal states (Terrazas et al., 2005), or sets of behavioral contingencies (Markus et al., 1995; D. M. Smith & Mizumori, 2006; Wood, Dudchenko, Robitsek, & Eichenbaum, 2000). One account of these findings holds that the hippocampus generates a "context code," defined in terms of the instantaneous activity of the hippocampal neuronal population, that becomes associated with the contents (e.g., stimuli and events) of the context (for a review see D. M. Smith & Bulkin, 2014). Consistent with this, rats are more likely to express contextual associations that match the current hippocampal context code (Bulkin, Law, & Smith, 2016; Kelemen & Fenton, 2010), and artificial reactivation of the hippocampal context code causes mice to recall context memories (Liu et al., 2012; Ramirez et al., 2013).

A common conjecture about the hippocampus is that its importance for memory processing is largely due to its position as the nexus of neocortical input. Indeed, many

neocortical regions also contribute to context processing, including the postrhinal cortex (Bucci, Phillips, & Burwell, 2000; Burwell, Bucci, Sanborn, & Jutras, 2004), the medial entorhinal cortex (Ferbinteanu, Holsinger, & McDonald, 1999), the medial prefrontal cortex (Jo et al., 2007), and the RSC (Keene & Bucci, 2008a, 2008b). And some even show hippocampal-like shifts in their neural activity when rats move between contexts, such as the medial entorhinal cortex (Fyhn, Hafting, Treves, Moser, & Moser, 2007), medial prefrontal cortex (Hyman, Ma, Balaguer-Ballester, Durstewitz, & Seamans, 2012), and the RSC, as described in chapter 2. In one popular model of hippocampal-neocortical interactions in memory processing, the hippocampus receives sensory input from cortical afferents and then drives recall by reactivating relevant representations stored in the neocortex (Marr, 1971; Miller, Vedder, Law, & Smith, 2014; Nadel, Hupbach, Gomez, & Newman-Smith, 2012; Teyler & Rudy, 2007). In support of this, artificially reactivating the hippocampal context code also reactivates context-specific representations in the neocortex (Tanaka et al., 2014). Furthermore, direct reactivation of neocortical LTM regions, such as the RSC, can result in recall even when the hippocampus is inactivated (Cowansage et al., 2014), suggesting that the primary role of the hippocampus in recall is neocortical reactivation.

Self-organizing Memory Consolidation

Learning and other cognitive tasks require continuously integrating new experiences into pre-existing memory representations. A vast literature investigating the biological changes occurring around the time of new learning has revealed two interdependent consolidation processes: an early, local cellular process that critically depends on hippocampal LTP, and a late-stage systems process whereby different structures show time-dependent roles in long-term

memory maintenance. However, recent research into how existing LTM affects new learning has shown that well-trained animals can often bypass this traditional consolidation timeline by simply updating the memories that they already have (i.e. reconsolidation; Nader & Hardt, 2009; 2000). During reconsolidation, fully consolidated memories are rendered labile again by a retrieval event, thereby allowing for the accommodation of novel information into the existing memory trace (Dudai, 2004).

Some of the most exiting data related to memory updating come from studies of rapid learning in rats with neocortical schemas (Tse et al., 2007; Tse et al., 2011). In this work, animals are cued with flavored treats to find rewards at spatial locations in an event arena. Training anew on this difficult task normally requires 4 – 6 weeks. Remarkably, however, fully trained animals can learn a completely new flavor-location pair in a single trial, and they consolidate this memory out of the hippocampus and into the neocortex (typically a weeks-long process) in less than 48 hours. Consistent with the idea that LTM supports this rapid learning, neocortical regions, including the RSC, show increased IEG expression only when rats rapidly acquire the new pair. Similar memory updating processes may also support incremental learning over many repeated exposures (J. L. Lee, 2008).

In addition to providing a mechanism by which LTM supports new learning, this research suggests that new memories may "self organize" into relevant LTM structures, with new information more easily learned when it is relevant to, and can therefore be accommodated into, existing LTM. Consistent with this, some recent work has suggested that the ability of novel information to be integrated into existing memories may depend crucially upon the match between the novel information and the existing LTM. This is because existing memories must be rendered labile by a destabilization process before novel information can be accommodated (S.-

H. Lee et al., 2012; S. H. Lee et al., 2008), and this destabilization process is affected by the congruency (van Kesteren, Ruiter, Fernandez, & Henson, 2012) between a stimulus and existing LTM associations. Thus, destabilization (and therefore reconsolidation) is driven by the presence of novelty, defined as incongruence between a stimulus and the expectations (Milekic & Alberini, 2002; Sevenster, Beckers, & Kindt, 2013; Suzuki et al., 2004).

Furthermore, it is important to remember that destabilization, when it does occur, does not affect all memories currently stored by the subject. Instead, only the memory that is most "dominant" at the time of manipulation is affected. For example, in experiments where there is competition in memory between retrieving the target memory and forming a new memory (Bouton, 2004), reconsolidation manipulations will selectively disrupt the memory that is currently being expressed by the rat (Eisenberg, Kobilo, Berman, & Dudai, 2003; Suzuki et al., 2004). Similarly, in some cases reconsolidation will not occur outside of the original learning context (DeVietti & Holliday, 1972; Misanin, Miller, & Lewis, 1968; Tse et al., 2007), consistent with the role of environmental contexts in reactivating memory traces.

Concluding Remarks

The results presented in this dissertation support the idea that neocortical regions such as the RSC support cognitive processes through the reactivation of LTM representations. The learning and memory field has long been concerned with how the hippocampus supports new learning and how hippocampal activity correlates with cognitive states. However, recent insights have pointed to a crucial role for the neocortex in these processes, and argue that a complete understanding of memory and cognition will require additional examination of hippocampal-neocortical interactions. In particular, the finding that the RSC contains rich spatial and

contextual representations that develop with learning, and that these LTM representations come to flexibly represent the rat's future behaviors, argues that neocortical regions like the RSC should feature prominently in any description of spatial cognition. Furthermore, there is a growing recognition that learning and memory are both continuous and dynamic: employing memories about the past to make predictions about the future, and, in turn, using predictions about the future to better form new memories. Teasing apart the contributions of various brain regions to this process will require new methods for simultaneously examining regional representations over long periods, as well as creative behavioral techniques for studying LTM during learning.

References

Anagnostaras, S. G., Gale, G. D., & Fanselow, M. S. (2001). Hippocampus and contextual fear conditioning: recent controversies and advances. *Hippocampus, 11*(1), 8-17. doi:10.1002/1098-1063(2001)11:1<8::aid-hipo1015>3.0.co;2-7

Barlow, H. (1961). Possible principles underlying the transformation of sensory messages. *Sensory Communication*, 217-234. doi:citeulike-article-id:7103098

Barlow, H. (2001). Redundancy reduction revisited. *Network, 12*(3), 241-253.

Bouton, M. E. (2004). Context and behavioral processes in extinction. *Learn Mem, 11*(5), 485-494. doi:11/5/485 [pii]

Bransford, J. D., & Johnson, M. K. (1972). Contextual prerequisites for understanding: Some investigations of comprehension and recall. *Journal of Verbal Learning & Verbal Behavior, 11*(6), 717-726. doi:10.1016/S0022-5371(72)80006-9

Brewer, W. F., & Treyens, J. C. (1981). Role of schemata in memory for places. *Cogn Psychol, 13*(2), 207-230. doi:10.1016/0010-0285(81)90008-6

Bucci, D. J., Phillips, R. G., & Burwell, R. D. (2000). Contributions of postrhinal and perirhinal cortex to contextual information processing. *Behav Neurosci, 114*(5), 882-894.

Bulkin, D. A., Law, L. M., & Smith, D. M. (2016). Placing memories in context: Hippocampal representations promote retrieval of appropriate memories. *Hippocampus, 26*(7), 958-971. doi:10.1002/hipo.22579

Burwell, R. D., Bucci, D. J., Sanborn, M. R., & Jutras, M. J. (2004). Perirhinal and postrhinal contributions to remote memory for context. *J Neurosci, 24*(49), 11023-11028. doi:10.1523/jneurosci.3781-04.2004

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cogn Psychol, 4*(1), 55-81. doi:10.1016/0010-0285(73)90004-2

Cowansage, K. K., Shuman, T., Dillingham, B. C., Chang, A., Golshani, P., & Mayford, M. (2014). Direct reactivation of a coherent neocortical memory of context. *Neuron, 84*(2), 432-441. doi:10.1016/j.neuron.2014.09.022

Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *J Exp Psychol, 58*(1), 17-22.

DeVietti, T. L., & Holliday, J. H. (1972). Retrograde amnesia produced by electroconvulsive shock after reactivation of a consolidated memory trace: A replication. *Psychonomic Science, 29*(3), 137-138.

Dudai, Y. (2004). The neurobiology of consolidations, or, how stable is the engram? *Annu Rev Psychol, 55*, 51-86. doi:10.1146/annurev.psych.55.090902.142050 [doi]

Eisenberg, M., Kobilo, T., Berman, D. E., & Dudai, Y. (2003). Stability of retrieved memory: inverse correlation with trace dominance. *Science, 301*(5636), 1102-1104. doi:10.1126/science.1086881 [doi]

Fenske, M. J., Aminoff, E., Gronau, N., & Bar, M. (2006). Top-down facilitation of visual object recognition: object-based and context-based contributions. *Prog Brain Res, 155*, 3-21. doi:10.1016/s0079-6123(06)55001-0

Ferbinteanu, J., Holsinger, R. M., & McDonald, R. J. (1999). Lesions of the medial or lateral perforant path have different effects on hippocampal contributions to place learning and on fear conditioning to context. *Behav Brain Res, 101*(1), 65-84.

Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *J Opt Soc Am A, 4*(12), 2379-2394.

Fyhn, M., Hafting, T., Treves, A., Moser, M.-B., & Moser, E. I. (2007). Hippocampal remapping and grid realignment in entorhinal cortex [10.1038/nature05601]. *Nature, 446*(7132), 190-194.

Godden, D., & Baddeley, A. (1975). Context-dependent memory in two natural environments [Article]. *British journal of psychology, 66*, 325-331.

Holland, P. C., & Bouton, M. E. (1999). Hippocampus and context in classical conditioning. *Curr Opin Neurobiol, 9*(2), 195-202.

Hyman, J. M., Ma, L., Balaguer-Ballester, E., Durstewitz, D., & Seamans, J. K. (2012). Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proc Natl Acad Sci U S A, 109*(13), 5086-5091. doi:10.1073/pnas.1114415109

Jo, Y. S., Park, E. H., Kim, I. H., Park, S. K., Kim, H., Kim, H. T., & Choi, J. S. (2007). The medial prefrontal cortex is involved in spatial memory retrieval under partial-cue conditions. *J Neurosci, 27*(49), 13567-13578. doi:10.1523/jneurosci.3589-07.2007

Keene, C. S., & Bucci, D. J. (2008a). Contributions of the retrosplenial and posterior parietal cortices to cue-specific and contextual fear conditioning. *Behav Neurosci, 122*(1), 89-97. doi:10.1037/0735-7044.122.1.89

Keene, C. S., & Bucci, D. J. (2008b). Neurotoxic lesions of retrosplenial cortex disrupt signaled and unsignaled contextual fear conditioning. *Behav Neurosci, 122*(5), 1070-1077. doi:10.1037/a0012895

Kelemen, E., & Fenton, A. A. (2010). Dynamic grouping of hippocampal neural activity during cognitive control of two spatial frames. *PLoS Biol, 8*(6), e1000403. doi:10.1371/journal.pbio.1000403

Komorowski, R. W., Manns, J. R., & Eichenbaum, H. (2009). Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. *J Neurosci, 29*(31), 9918-9929. doi:10.1523/jneurosci.1378-09.2009

Lee, J. L. (2008). Memory reconsolidation mediates the strengthening of memories by additional learning. *Nat Neurosci, 11*(11), 1264-1266. doi:nn.2205 [pii]

Lee, S.-H., Kwak, C., Shim, J., Kim, J.-E., Choi, S.-L., Kim, H. F., . . . Kaang, B.-K. (2012). A cellular model of memory reconsolidation involves reactivation-induced destabilization and restabilization at the sensorimotor synapse in Aplysia. *Proceedings of the National Academy of Sciences, 109*(35), 14200-14205. doi:10.1073/pnas.1211997109

Lee, S. H., Choi, J. H., Lee, N., Lee, H. R., Kim, J. I., Yu, N. K., . . . Kaang, B. K. (2008). Synaptic protein degradation underlies destabilization of retrieved fear memory. *Science, 319*(5867), 1253-1256. doi:1150541 [pii]

Liu, X., Ramirez, S., Pang, P. T., Puryear, C. B., Govindarajan, A., Deisseroth, K., & Tonegawa, S. (2012). Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature, 484*(7394), 381-385. doi:10.1038/nature11028

Markus, E. J., Qin, Y. L., Leonard, B., Skaggs, W. E., McNaughton, B. L., & Barnes, C. A. (1995). Interactions between location and task affect the spatial and directional firing of hippocampal neurons. *J Neurosci, 15*(11), 7079-7094.

Marr, D. (1971). Simple Memory: A Theory for Archicortex [10.1098/rstb.1971.0078]. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 262*(841), 23-81.

Milekic, M. H., & Alberini, C. M. (2002). Temporally graded requirement for protein synthesis following memory reactivation. *Neuron, 36*(3), 521-525. doi:S0896627302009765 [pii]

Miller, A. M. P., Vedder, L. C., Law, L. M., & Smith, D. M. (2014). Cues, context, and long-term memory: the role of the retrosplenial cortex in spatial cognition [Review]. *Frontiers in Human Neuroscience, 8*. doi:10.3389/fnhum.2014.00586

Misanin, J. R., Miller, R. R., & Lewis, D. J. (1968). Retrograde amnesia produced by electroconvulsive shock after reactivation of a consolidated memory trace. *Science, 160*(3827), 554-555.

Morris, R. G., Garrud, P., Rawlins, J. N., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature, 297*(5868), 681-683.

Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J Neurosci, 7*(7), 1951-1968.

Nadel, L., Hupbach, A., Gomez, R., & Newman-Smith, K. (2012). Memory formation, consolidation and transformation. *Neurosci Biobehav Rev, 36*(7), 1640-1645. doi:10.1016/j.neubiorev.2012.03.001

Nader, K., & Hardt, O. (2009). A single standard for memory: the case for reconsolidation. *Nat Rev Neurosci, 10*(3), 224-234. doi:nrn2590 [pii]

Nader, K., Schafe, G. E., & Le Doux, J. E. (2000). Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval. *Nature, 406*(6797), 722-726. doi:10.1038/35021052 [doi]

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res, 34*(1), 171-175.

Ramirez, S., Liu, X., Lin, P. A., Suh, J., Pignatelli, M., Redondo, R. L., . . . Tonegawa, S. (2013). Creating a false memory in the hippocampus. *Science, 341*(6144), 387-391. doi:10.1126/science.1239073

Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*(4), 803-814. doi:10.1037/0278-7393.21.4.803

Sevenster, D., Beckers, T., & Kindt, M. (2013). Prediction error governs pharmacologically induced amnesia for learned fear. *Science, 339*(6121), 830-833. doi:339/6121/830 [pii]

Smith, D. M., & Bulkin, D. A. (2014). The form and function of hippocampal context representations. *Neurosci Biobehav Rev, 40c*, 52-61. doi:10.1016/j.neubiorev.2014.01.005

Smith, D. M., & Mizumori, S. J. (2006). Learning-related development of context-specific neuronal responses to places and events: the hippocampal role in context processing. *J Neurosci, 26*(12), 3154-3163. doi:10.1523/jneurosci.3234-05.2006

Smith, S., Glenberg, A., & Bjork, R. (1978). Environmental context and human memory. *Memory & Cognition, 6*(4), 342-353. doi:citeulike-article-id:9245897

Smith, S. M., & Vela, E. (2001). Environmental context-dependent memory: a review and meta-analysis. *Psychon Bull Rev, 8*(2), 203-220.

Suzuki, A., Josselyn, S. A., Frankland, P. W., Masushige, S., Silva, A. J., & Kida, S. (2004). Memory reconsolidation and extinction have distinct temporal and biochemical signatures. *J Neurosci, 24*(20), 4787-4795. doi:10.1523/JNEUROSCI.5491-03.2004 [doi]

Talk, A. C., Gandhi, C. C., & Matzel, L. D. (2002). Hippocampal function during behaviorally silent associative learning: dissociation of memory storage and expression. *Hippocampus, 12*(5), 648-656. doi:10.1002/hipo.10098

Tanaka, Kazumasa Z., Pevzner, A., Hamidi, Anahita B., Nakazawa, Y., Graham, J., & Wiltgen, Brian J. (2014). Cortical Representations Are Reinstated by the Hippocampus during Memory Retrieval. *Neuron, 84*(2), 347-354. doi:10.1016/j.neuron.2014.09.037

Terrazas, A., Krause, M., Lipa, P., Gothard, K. M., Barnes, C. A., & McNaughton, B. L. (2005). Self-motion and the hippocampal spatial metric. *J Neurosci, 25*(35), 8085-8096. doi:10.1523/jneurosci.0693-05.2005

Teyler, T. J., & Rudy, J. W. (2007). The hippocampal indexing theory and episodic memory: updating the index. *Hippocampus, 17*(12), 1158-1169. doi:10.1002/hipo.20350

Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., . . . Morris, R. G. (2007). Schemas and memory consolidation. *Science, 316*(5821), 76-82. doi:10.1126/science.1135935

Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., . . . Morris, R. G. (2011). Schema-dependent gene activation and memory encoding in neocortex. *Science, 333*(6044), 891-895. doi:10.1126/science.1205274

van Kesteren, M. T., Ruiter, D. J., Fernandez, G., & Henson, R. N. (2012). How schema and novelty augment memory formation. *Trends Neurosci, 35*(4), 211-219. doi:10.1016/j.tins.2012.02.001

Wood, E. R., Dudchenko, P. A., Robitsek, R. J., & Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron, 27*(3), 623-633.

Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *Trends Cogn Sci, 10*(7), 301-308. doi:10.1016/j.tics.2006.05.002