ORGANIZATION OF SPATIOTEMPORAL INFORMATION AND RELATIONAL MEMORY IN THE HIPPOCAMPUS

BY

KEVIN MICHAEL HORECKA

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Neuroscience
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2018

Urbana, Illinois

Doctoral Committee:

Professor Neal J. Cohen, Chair
Dr. Rama Ratnam
Associate Professor Aron K. Barbey
Associate Professor Arne Ekstrom, University of Arizona

ABSTRACT

This work examines the role of the hippocampus and relational memory in organizing episodic memory during navigation and reconstruction. Navigation is a critical component in most organisms' survival. Reconstruction, on the other hand, provides an incredibly rich method of evaluating the precise information remembered by an individual after attempting to learn and remember that information. Through validating the computational framework in this work on amnesic patients with hippocampal damage, an understanding of some of the specific *types* of relations which rely on the hippocampus can be established. Then, this framework can be applied to a much more complex, spatiotemporal navigation and reconstruction task in healthy individuals to gain a wider perspective on the organization of episodic memory, which is known to critically rely on the hippocampus.

The first experiment and associated analysis framework presented in this document (Chapter 2) uses spatial reconstruction to establish that not all *types* of spatial relations are impaired in hippocampal damaged patients. In particular, the arbitrary, identity-location relations (i.e. those relationships where the element being bound could have just as easily been anything) are critically impaired in hippocampal damaged patients while location information, disregarding identity, is not. The use of reconstruction in this context allows for the establishment of a set of critical computational metrics which relate to hippocampal function in reconstruction which can then be applied to other reconstruction tasks in healthy individuals to learn more about the wider structure and organization of memory.

In the second experiment (Chapters 3 and 4), the methodologies which were applied to hippocampal damaged patients in the first experiment are applied to a novel Spatiotemporal Navigation Task in healthy young adults. In this task, participants are not just asked to study and reconstruct items in space, but instead, participants are asked to, in Virtual Reality, navigate space and time (via normal movement and simulated Time Travel) and study, then reconstruct the locations of events in spacetime. The computational framework established in the previous chapter is then applied to show that relational memory errors in time are far more common in this task than in space, suggesting differences in representations between these two domains even when the navigation and exploration of the domains are put on a more equal footing. Additionally, in time, these relational memory errors are far more likely to occur within a shared contextual region than should occur by chance. In fact, this error (temporal relational memory error within a context) gets worse across the first 3 trials, suggesting a systematic bias due to context. Finally, a more traditional bias, the context boundary effect (i.e. a "squishing" of

ABSTRACT

This work examines the role of the hippocampus and relational memory in organizing episodic memory during navigation and reconstruction. Navigation is a critical component in most organisms' survival. Reconstruction, on the other hand, provides an incredibly rich method of evaluating the precise information remembered by an individual after attempting to learn and remember that information. Through validating the computational framework in this work on amnesic patients with hippocampal damage, an understanding of some of the specific *types* of relations which rely on the hippocampus can be established. Then, this framework can be applied to a much more complex, spatiotemporal navigation and reconstruction task in healthy individuals to gain a wider perspective on the organization of episodic memory, which is known to critically rely on the hippocampus.

The first experiment and associated analysis framework presented in this document (Chapter 2) uses spatial reconstruction to establish that not all *types* of spatial relations are impaired in hippocampal damaged patients. In particular, the arbitrary, identity-location relations (i.e. those relationships where the element being bound could have just as easily been anything) are critically impaired in hippocampal damaged patients while location information, disregarding identity, is not. The use of reconstruction in this context allows for the establishment of a set of critical computational metrics which relate to hippocampal function in reconstruction which can then be applied to other reconstruction tasks in healthy individuals to learn more about the wider structure and organization of memory.

In the second experiment (Chapters 3 and 4), the methodologies which were applied to hippocampal damaged patients in the first experiment are applied to a novel Spatiotemporal Navigation Task in healthy young adults. In this task, participants are not just asked to study and reconstruct items in space, but instead, participants are asked to, in Virtual Reality, navigate space and time (via normal movement and simulated Time Travel) and study, then reconstruct the locations of events in spacetime. The computational framework established in the previous chapter is then applied to show that relational memory errors in time are far more common in this task than in space, suggesting differences in representations between these two domains even when the navigation and exploration of the domains are put on a more equal footing. Additionally, in time, these relational memory errors are far more likely to occur within a shared contextual region than should occur by chance. In fact, this error (temporal relational memory error within a context) gets worse across the first 3 trials, suggesting a systematic bias due to context. Finally, a more traditional bias, the context boundary effect (i.e. a "squishing" of

within context temporal locations and "stretching" of across context temporal locations) is observed even though participants are allowed to reexplore the contexts arbitrarily, multiple times. This suggests that the context boundaries are having a profound impact on both the distance judgements and relational memory structure associated with events in spacetime.

Finally, in the fourth chapter, the navigation component of the previous Spatiotemporal Navigation Task is examined to determine if changes in study time navigation and exploration relate to changes in the various test metrics discussed in the previous chapter. More rapid improvements in spatial and temporal navigation are shown to relate to more rapid improvements in memory in those domains, separably, suggesting that spatial and temporal representations may in some way be separable in this task in both the relational representations and the navigation strategies supporting those representations. Relational memory improvements are shown to be uniquely tied to changes in navigation complexity and systematicity, pointing to an interplay between in-the-moment, memory-guided decision making and subsequent relational memory efficacy. Context boundaries are suggested to act as more of a discriminatory feature (at least in this task) than one used to strengthen within-context relational memory organization accuracy as there is a significant relationship between changes in context boundary crossing and both the context boundary effect and across-context temporal relational memory errors. Finally, a preference towards exploring an otherwise temporally-flexible environment in the implied, forward order with increasing contiguity is suggested to be a critical element in improving temporal, relational, and contextual memory organization.

Taken together, this work shows the richness of spatiotemporal navigation and reconstruction in observing the complex interplay between navigation in space, navigation in time and how these ultimately may relate to navigation in memory. Through embracing principled approaches to analysis of behavioral data, and the inclusion of complex behavioral mechanics (such as simulated time travel), this work extends our understanding of the role of hippocampal relational memory and overall memory organization.

ACKNOWLEDGEMENTS

Of course, my time in Illinois has not just been research, and maintaining a reasonable mental health while completing a PhD is no small task. The friends I have made here are ones I will always cherish. You all made being here more than just a job – you made it a life. To Elizabeth Davis, Sean Collins, Angie Bustamante, Carly Drzewiecki, Jean Carlos Morganti, Francheska Marie Morganti Nievesm, Lydia Nguyen, Brett Velez, and Stephen Fleming, thank you. And to all the friends who let us stay with them on our dozens of visits home to Austin, TX, Kevin Behlmann, Chris Martin, Divya Janardhana, Brette Hannigan, Eric and Sarah Hofstetter, I love you guys. To my family, Cathy and Michael Horecka, Katie Bares and the entire Bares family, and my grandparents, Joan Bauer, Bill Bauer, Paul Horecka, and LaVonne Horecka, thank you for everything you've done for me throughout my life – I wouldn't be here without you.

Finally, to my wife, Nipuni Ratnayaka – I'm so happy we were able to come to Illinois together and both achieve our dreams. I love you so much, and I am so excited to continue to the next phase of our lives together as doctors. Also, special thanks to Timber and Sashi, for their patience with the cold.

To my wife, friends, and family

And

All of my mentors and teachers

Contents

Chapter 1: General Introduction

There are numerous ways to quantify memory efficacy, but when it comes to recollective memory few methods are richer or more detailed than navigation (Eichenbaum, 2017b) and reconstruction tasks (Huttenlocher & Presson, 1979). Reconstruction tasks ask individuals to study some set of information and then, after a delay, recreate the information they observed. Traditionally, individuals are given a set of items which are to be reconstructed in space (Jeneson et al., 2010; Smith & Milner, 1981; Watson et al., 2013). The strength of these paradigms is that not only can one assess which aspects of the original information were or weren't remembered, but also *the degree to which* they were remembered, as well as whether there are any systematic distortions in that memory. Ultimately, we are unable to know what precise thoughts, strategies, or information the participant did or did not have when performing the task just from the behavioral data, but by carefully picking the assumptions we are willing to make, we can strongly infer what was not done in the reconstruction (because if it were, the performance should have been different). Moreover, by systematically analyzing the performance in reconstruction, we can begin to hypothesize new models and test existing models of memory organization (Eichenbaum, 2016) and representation (Eichenbaum, 2004, 2017a; O'Reilly & Rudy, 2001; Schapiro et al., 2017) across many different domains of information (Konkel et al., 2008).

In addition to a theoretical introduction to assumptions and perspectives on reconstruction tasks, the work presented in this document will show how performance on reconstruction tasks can be broken down into specific and sensitive metrics of performance to elucidate some aspects of memory in two particular domains, space and time. Relational memory theory posits that the hippocampus is critical for the binding of arbitrary relations into rich memory representations, as well as for the flexible usage of these representations, regardless of the domain from which the information originates (Cohen & Eichenbaum, 1993; Monti et al., 2015). Previous work has shown that binding of the identity of an item to a location in space (i.e. identity-location binding) may be one critical example of an arbitrary relation which is severely impaired in hippocampal damaged patients (Watson et al., 2013), with other accounts suggesting that not all relational information is impaired (Jeneson et al., 2010). Though these previous accounts hinted at specific aspects of information in reconstruction being impaired and spared, the analysis methods differed significantly, making it difficult to determine the precise nature of the deficits in reconstruction in hippocampal damaged patients. The work presented here attempts to reconcile these accounts with a novel methodological framework constructed from a principled approach to reconstruction data. The data show that certain (but not all) *types* of relational information are impaired

in hippocampal damaged individuals (Chapter 2), but, critically, arbitrary identity-location information is impaired (i.e. the information was not derivable from other relations or elements which may have been remembered) while the location information (i.e. the ability to reconstruct items in studied locations regardless of the identity of the item) is not.

Hippocampal function is not exclusive to spatial information, however, with the identification of "time cells" in the hippocampus recently showing this via direct, electrophysiological evidence (Eichenbaum, 2014; Kraus et al., 2013). Additional sources of information (such as contextual information) can also bias representations (Ezzyat & Davachi, 2014; Zacks et al., 2007), and these additional sources of information are critical in understanding the overall organization of memory and the hippocampal system, with most modern models of the region including cortical region interactions which involve the addition of these alternate sources of information (Kumaran et al., 2016; Sederberg et al., 2008; Wang et al., 2015). The reconstruction task presented in Chapter 3 embrace this complexity by allowing space and time to be explored simultaneously and asking participants to reconstruct events in space-time in the presence of temporal contexts. The data from this task show that in healthy adults, there may be differences in the representations of spatial and temporal information and systematic biases in different aspects of temporal representations due to contextual information.

Finally, sampling behavior during study in reconstruction tasks is an equally critical element in overall performance. In previous work involving spatial sampling, eye movements, which are often thought of as "visual exploration" of an environment/scene (Voss et al., 2017), in hippocampal damaged individuals and healthy individuals indicated both memory for an entire previously viewed scene and relations among elements in the scene; however, when the relations between items were changed, only healthy individuals showed differences in eye movements relative to the change, while hippocampal damaged patients did not (Cohen & Banich, 2003). Similar sampling deficits in spatial navigation have been shown in rodents with hippocampal damage in the Morris Water Maze (Morris et al., 1982). However, more recent work in humans shows that this deficit in navigation may be exclusive to precise, high-resolution information as humans with hippocampal damage in a Virtual Morris Water Maze were able to navigate to the correct quadrant of the spatial area but were impaired in finding the precise location of the target (Kolarik et al., 2016; Kolarik et al., 2017). The hippocampus' role in spatial navigation has been compared extensively across species (especially rodents and humans), and the subject is still hotly debated (Eichenbaum, 2017c). However, as stated in the previous paragraphs, hippocampal function is not exclusive to spatial processing and is critical for all domains of information. As such, this work (Chapter

4) examines exploration and navigation in a domain which has never been examined before, the temporal domain. Temporal navigation (aka Time Travel) may seem like a farfetched idea, but as many other authors all the way back to Tulving, 1989 (and more explicitly in Tulving, 2002a) have said, memory is time travel, with further evidence that MTL neurons "jump back in time" by replaying specific patterns of firing in response to repeated stimulus (Howard et al., 2012). The data in this chapter show that certain aspects of learning in navigation and exploration of both time and space relate significantly to learning of relational and contextual information, and ultimately, the systematicity and complexity of navigation may be one of the more critical components in determining which individuals will improve most significantly in their temporal relational memory.
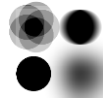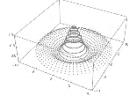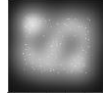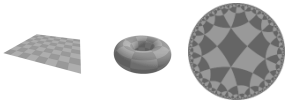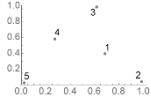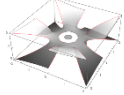
Domains	Entities
Finite vs. Infinite (i.e. there are boundaries or not) 	**Precisely Located** and **Encapsulated/Finite** (i.e. their position can be defined by one point)
Continuous vs. Discrete (i.e. any location is valid or only some are) 	**Distribution** (i.e. entities are more likely in some areas)
Dimensionality (i.e. there are multiple of the same domain type) 	**Constraint** (i.e. positions of entities are relative to others)
Geometry/Topology (i.e. Euclidean, toroidal, hyperbolic) 	**Identity** (i.e. they are anonymous or labelled)
Disjoint vs. Connected (i.e. can all space be reached from anywhere) 	**Category** (i.e. entities have distinct groupings)

Table 1.1: Properties of Domains and Entities. This list is not necessarily exhaustive, but these properties represent some critical assumptions being made during any reconstruction task.

Before examining specific data on reconstruction and hippocampal function, several details regarding the theoretical prerequisites to analysis of this data should be explored. Information in reconstruction can be organized in a huge variety of ways, and it is useful to ensure that the analysis methods are capable of handling as many of these cases as possible. Also, by examining the assumptions made during task design, we can not only gain better insights into the nature of the task and analyses but we are also presented with an enormous variety of follow-up experiments and questions that future work outside of this dissertation may address. The next section breaks down properties of this task by examining the two critical elements of reconstruction, *domains* and *entities*, and enumerating properties and assumptions which warrant consideration in designing reconstruction tasks and analyses.

1.1 Domains and Entities: Building Systematic Understanding of Reconstruction

Reconstruction can be defined, in its most basic form, as containing two main components: a *domain,* i.e. an axis or set of axes into which *entities* (elements of a domain such as items, events, precise colors, pitches, etc.) can be embedded (see **Table 1.1** for an enumeration of key properties of domains and entities). It is interesting to see the similarities between this classification and the dorsal and ventral "where" and "what" pathways of the perceptual system (Ungerleider & Mishkin, 1982), which converge in the hippocampal system (Insausti et al., 1987; Squire & Zola-Morgan, 1991; Van Hoesen et al., 1979) as well as in descriptions of processing in scene-selective visual areas such as the functionally defined Parahippocampal place area (see Epstein, 2014 for a detailed examination). While, of course, multiple types of domains and entities interact in memory, for the sake of our understanding of the memory representations which are involved in reconstruction, it is helpful to at first restrict our focus to instances in which only one domain and one type of entity are present, at least when developing our methodological and theoretical framework.

Figure 1.1. The Beckman Institute; an example of a 3D spatial domain populated by numerous entities which occupy specific points in space and have various identities. Red/dotted squares highlight entities which occupy the spatial domain. Some occupy fixed locations while others move at different time scales. The blue/dotted arrows illustrate the three spatial dimensions which comprise the spatial domain.

Let us take a specific example to which we can link the abstract terms defined in the previous paragraph to clarify their meaning. Imagine a 2D spatial domain in a space such as the Beckman Institute Café (pictured in **Figure 1.1**). Within this spatial domain exist items, i.e. discrete entities which have a precisely defined location (such as the lamp, humans, and kiosk) and locations relative to one-another (as defined by the three spatial axes). If we consider the number of restrictions we have already made, we can begin to understand how difficult the problem of a systematic breakdown of reconstruction is. First, the reader may have implicitly assumed (likely, rightfully so due to my example) that the 2D space in question is finite. However, this was never explicitly stated (i.e. the spaces adjacent to the café, outside of the Beckman Institute, outside of Urbana, IL, and outside of Earth's atmosphere may be equally considered part of the space). An infinite space opens up a variety of new problems of scale and geometry which future studies should address. For example, space could be a toroidal ring like we see in video games such as the classic, Asteroids (Wolf, 2010), wrapped on a sphere like we experience on maps from day-to-day (i.e. the earth; Snyder, 1987), in a hyperbolic space (*a la* a Penrose diagram in theoretical physics; Hawking & Ellis, 1973) where distance becomes asymptotically large as boundaries are approached, or something more exotic (Coxeter, 2008; see **Figure 1.2** for visualizations of some of these spaces). Although humans more typically navigate a flat, 2D, Euclidean space, we are nonetheless

capable of understanding and utilizing these other spaces (as can be easily observed by anyone who has ever played Asteroids or navigated the globe).
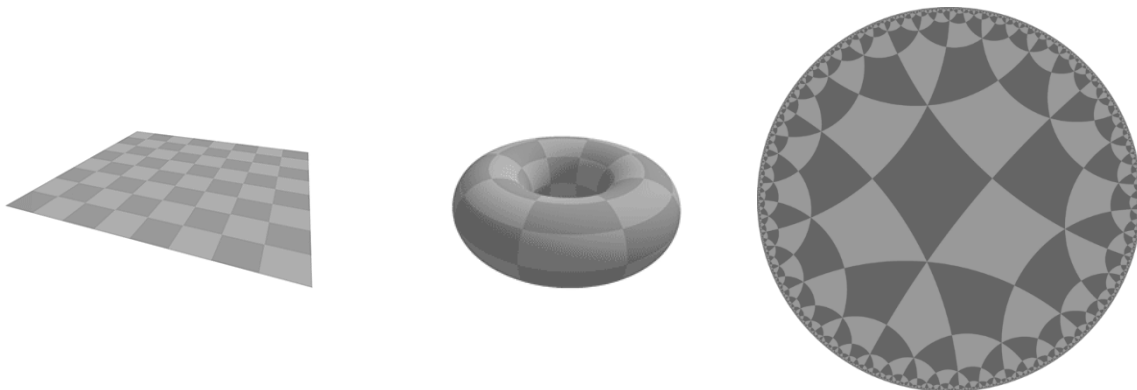

Figure 1.2: Examples of spaces with different geometries. Euclidean, toroidal, and hyperbolic.

Entities, on the other hand are typically restricted to have a precisely defined location, so entities which occupy multiple locations simultaneously have been eliminated. Entities typically have identities (rather than being anonymous) and are assumed to be equally likely to be positioned anywhere in the space while not occupying the same location as another entity (though we will see in the next section how subtle violations of this assumption can warp reconstruction expectations). These systematic assumptions that both experimenter and participant make may seem unimportant as, if both parties make the same assumptions, it should not confound the results. However, participants who are unable to remember some information due to an impairment such as hippocampal damage may fall back to these core assumptions and perform the task via heuristic, guessing locations based on an understanding of task design rather than an explicit knowledge of locations. Minor changes in how the studied locations are chosen can have a large impact on the subsequent probability space of the items and, therefore, drastically change the efficacy of various heuristics in chosen locations (see **Figure 1.3** for examples of various probably spaces from real reconstruction experiments). There is mounting evidence that hippocampal damage impairs precision in spatial memory (Kolarik et al., 2016, 2017; Yonelinas, 2013), which could be explained, in part, by a heuristic approach to memory. Alternatively, patients might actually be able to maintain relational information of some sort via alternative methods. Patients might, for instance, remember Gestalten perceptual features, i.e. configural features which constitute parts of a unified whole, rather than specific types of relations in order to reconstruct (Corsi, 1972; Kessels et al., 2000; Uttal & Chiong, 2004). Although we might be tempted to call heuristic performance "random" (because on some level, patients are "guessing" locations without any precise knowledge), or perceptually based performance "imprecise", by acknowledging when our own assumptions of our task

might become violated in a systematic way, we might gain insight into the sorts of information processing which are and are not impaired with damage to certain brain regions, and, by extension, we may better understand neural information processing as a whole.
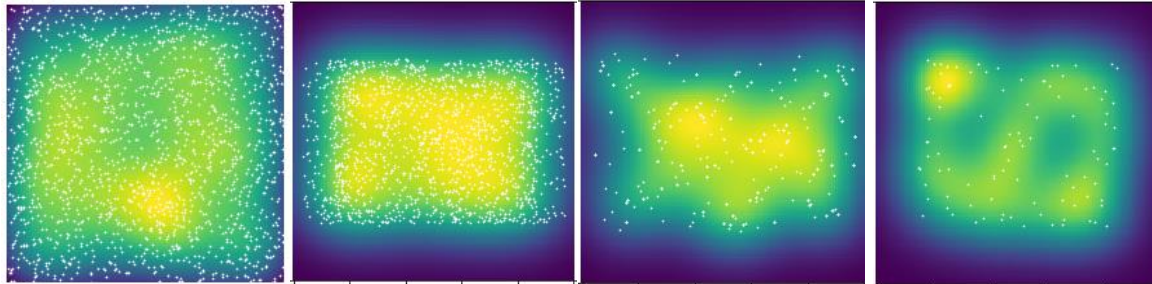

Figure 1.3: Examples of heatmaps of item placement with different constraints. From left to right: random placement, random placement with aspect ratio constraint, random placement avoiding boundaries, random placement with a distance constraint (i.e. items must be a certain distance apart).

1.1.1 Illustratability Does Not Define Domains of Information

Another critical assumption often made in reconstruction is the types of domains and entities which might be involved. Typically, reconstruction tasks have involved studying entities located in the spatial domain (as illustrated above). However, this is, once again, an arbitrary imposition on what reconstruction is, especially given the evidence that the hippocampus is critical for all domains of information (Konkel et al., 2008). Before enumerating examples of domains and entities which go beyond these examples, it's useful to take a moment to consider why space (and sometimes time) may be so dominant. Reality, as we observe it, is intrinsically experienced in 3 spatial dimensions and 1 temporal dimension. All information of any type which we observe will be embedded in these 4 dimensions. This is often used as an argument as to why these must be the critical dimensions to neural systems, and perhaps, on some level, they are. However, it is incredibly restrictive to allow the definitions of domains which could be encoded to only be those which can be illustrated. Illustration requires, by its very nature, embedding whatever domain or idea is being observed into the observable 4D world. However, the set of all possible representable information need not be bound to what can be illustrated. An analogous situation happens when we imagine the dimensionality of data. We cannot easily illustrate data beyond 3 dimensions (see Abbott, 1884 for a playful exploration of attempting to visualize alternate numbers of dimensions). We can add color, animation, shape, size, and all sorts of other illustrative methods, but we can always add more dimensions to the data. If we stopped accepting dimensions beyond those that can be illustrated as easily as we stop accepting domains of information,

much of the powerful mathematics used in the world today would be out of our reach. It is not a necessary precondition for a brain to represent information which can be illustrated.

1.1.2 Domains Other than Space and Time

Several domains other than space and time have been examined in the literature in recent years. Social space, as defined by the relative affiliation and power between individuals (Eichenbaum, 2015; Tavares et al., 2015), is another possible abstract space and has been investigated in the context of hippocampal function, finding that hippocampal fMRI activation located characters in a 2D power-affiliation "map". Color can be thought of as an abstract domain where entities can be embedded in locations in a "color space" independently of locations in space and time (see Warren et al., 2015 for an example of a reconstruction task involving color as a domain of interest). Of course, it is tempting to fall back into old habits and think of color as being on a 1D line, with the x axis defined as hue, and social space being illustrated via a scatter plot in 2 dimensions. These are useful illustrations, but domains of information are only illustrated to help us understand them. Their illustration does not define them. Additionally, although time has been discussed as if it can be easily lumped in with space up until now, it obviously has its own interesting properties. Time is intrinsically unidirectional. Time is continuous, and motion through it is obligatory. Few other domains have this property, and as such, time may be of special interest. Moreover time cells have been identified in the hippocampus (Kraus et al., 2013) which act much like place cells but activate corresponding with particular moments in time. Although it is difficult to disentangle temporal firing from spatial or distance firing, via careful task design, cells which fire to time and distance in exclusion of one another (as well as cells which fire for both) have been identified (Kraus et al., 2015). Together these pieces begin to paint a picture of a hippocampus in which entities can be bound within and across a variety of domains. Indeed, previous work has shown the hippocampus is critical for all manner (i.e. domain) of relations (Konkel et al., 2008). As such, our methodological framework should attempt to be agnostic to domain as much as is possible so that domain representation differences can be identified using the same metrics.

Relational memory theory (Cohen & Eichenbaum, 1993) captures the vast majority of phenomena described in the previous paragraphs, but several properties that have been described have, to my knowledge, never been tested. It is not the goal of this dissertation to systematically test all of these variations, but rather, by describing these variations, a better understanding of the simpler cases of reconstruction can be gained which can generalize to other domains with more complex properties. In this way, relational memory theory can be refined to determine the *types* of relationships which are

critical in different domains given different properties, and the organization of memory regardless of domain can become clearer. Although the complex cases described above of domains with toroidal geometry and disjoint features do exist in video games such as Asteroids or Portal and can be potentially addressed with the methods described in Chapter 2, a more fundamental understanding of reconstruction in the familiar domain of 2D Euclidean space should first be attained.

1.1.3 2D Spatial Reconstruction – The Spatial Reconstruction Task

In the parlance of the previous sections, the Spatial Reconstruction task is comprised of a finite, continuous, 2D, Euclidean, spatial, connected domain (i.e. a computer screen) populated with 6, precisely located, finite, randomly distributed pictures of real-world entities which will not overlap/appear on top of one another and are sampled from a finite set of possible entities. This scenario has been used in several publications at this point (Jeneson et al., 2010; Lucas et al., 2016; Monti et al., 2015; Schwarb et al., 2016, 2017; Watson et al., 2013) and has a proven track record of having hippocampal-dependent performance components. To break down this task, it is helpful to consider the information-theoretic content of the task. To determine the amount of information necessary to perfectly reconstruct a 6-item trial of the task, the simplest method is to assess how many numbers would be needed to describe 6 precise locations and the associated identity information. This question is intimately related to questions of information entropy (Cover, 2006; Shannon, 1948). A naïve approach to describing the location information would be to represent each of the 6 items with 2 numbers (x and y) and, assuming we have a list of the possible objects (or at least an understanding of the objects which are currently present), 1 additional number is needed to define the identity of each entity. So, in our naïve case, we can represent all 6 items with 18 numbers. However, as we'll see in a moment, this naïve representation may or may not represent the actual amount of information storage used by an individual in remembering the configuration in this task.
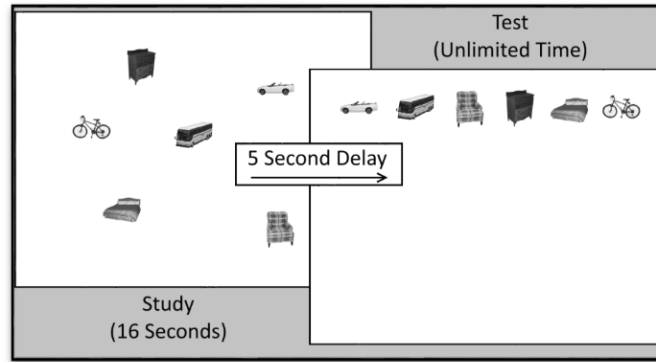
Figure 1.4: Example Reconstruction Task: a finite, continuous, 2D, Euclidean, spatial, connected domain populated with 6, precisely located, finite, randomly distributed pictures of real-world objects which will not overlap/appear on top of one another and are sampled from a finite set of possible objects. Note that the study portion is timed, there is a short delay, then the test has unlimited time. These are constraints on sampling the environment and maintaining memory more than the information contained within the configuration.

First, for location information, if the intent is to reconstruct an original configuration of these items, a precise reconstruction will require all 12 location numbers to be represented. However, for an approximate representation, a much more compressed option may be available. Suppose the positions of each element were encoded into a graph-like representation where the relations between a subset of the items was encoded instead of any explicit location. This is, in fact, a representational framework posited by relational memory theory (Cohen & Eichenbaum, 1993). Now, if the graph can be represented more efficiently, only one position need be remembered precisely (i.e. the overall position of the graph within the environment). Moreover, with this representation, the positions can now be transformed in ways which are not obvious given a set of 6 unassociated points. For instance, if the space were to be scaled up or down (as was done in Muller & Kubie, 1987 in rodents), the relative relations would simply be scaled up or down in accordance (and in fact, Muller & Kubie showed that place cell firing in an open field remapped with scaling, maintaining the relative shape of spatial firing). If the space were to be translated (i.e. moved in any direction a fixed distance), the graph itself need not change, only the global position would need to be adjusted. In fact, within a space like the one we've defined, scaling and translation are the fundamental transformations which do not require alteration of the graph of relative positions. Interestingly, some theories of temporal memory for order also share this property of scale and translation invariance in time (Howard et al., 2005; Howard & Kahana, 2002; Howard & Natu, 2005; Shankar et al., 2016). Rotation, sheering, or other more complex transformations can potentially result in the relative positions of items being changed, but with scaling and translation,

any item to the left of another will remain to the left, and any item on top will remain on top (see **Figure 1.5** for a visualization of these phenomena). This idea parallels similar ideas from physics and mathematics, where certain quantities are conserved (or invariant) under certain transformations. Noether's theorem (Noether, 1971), in fact, specifically states that for any transformation (specifically, those which are differential symmetries of an action), there is some corresponding conservation law (conservation of energy is equivalent to translational symmetry in time, conversation of moment is translational symmetry in space, conservation of angular momentum is rotational symmetry in space, etc.). If we believe memory for the location information which describes a set of points is encoded in such a relational graph, we should similarly believe that the representation should only be conserved under those transformations which have relational invariance (i.e. scaling and translation) but not those which lack such an invariance (i.e. rotation and sheering). Moreover, we might suspect that transformations which conserve parts of the relational information (like reflection, which conserves relations on all but the reflected axis) might have a graded effect on memory for the locations. However, in 1D domains (such as time) a reflection would have a more extreme effect as it is guaranteed to damage the relational information between every item pair that isn't directly on top of each other.
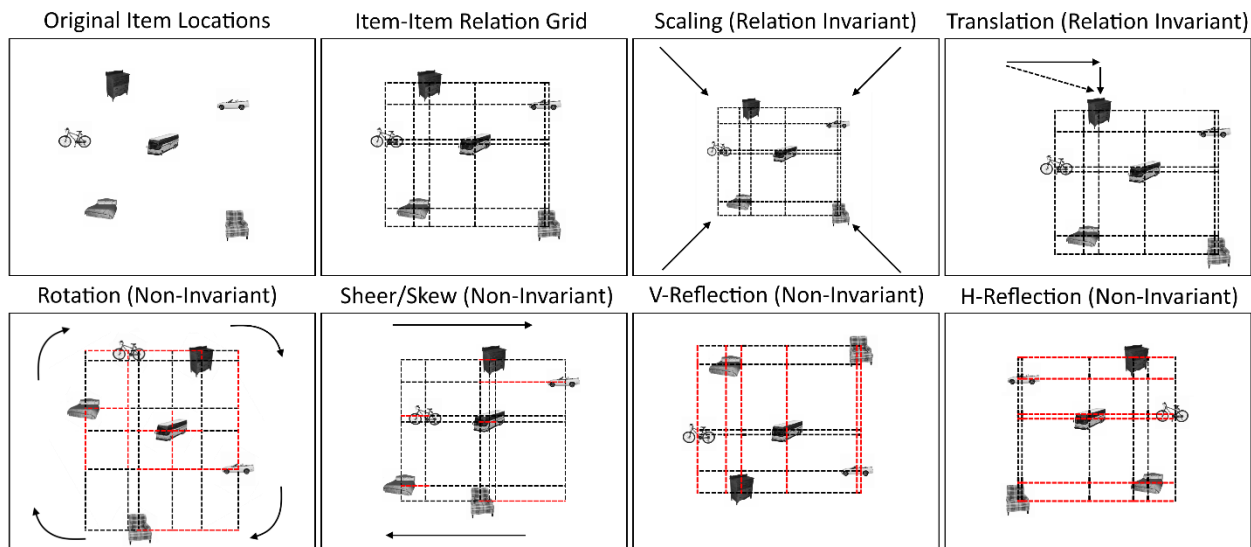
Figure 1.5: Relational Invariance in Affine Transformation Components in 2D: Note that via an item-item relational grid, we can observe the impact of different transforms on the relations present between items. Scaling and Translation have no impact on the relations between items (though scaling impacts distance between items. On the other hand, Rotation, Sheer/Skew, and Vertical/Horizontal Reflection all can modify relations in some systematic way. However, rotation uniquely can modify both dimensions relations.

Of course, the precision of the memory for the graph position and for the relational information may vary as well, but this can be treated as different types of noise which can be quantified in particular ways. The Precision and Binding Model (Yonelinas, 2013) suggests that the hippocampus is critical for binding precise, high resolution information, and the associated measures of precision used in testing this model (a variable sized accuracy window) are one example of a quantification method which shows promise in providing more interpretability when working with complex path data (Kolarik et al., 2016, 2017). However, this model addresses precision in continuous domains, not precision for discrete information, and it does not address the specifics of how different binding configurations might influence memory. Moreover, it is unclear how the concept of precision should be generalized for several stimuli being bound to each other as well as elements of the environment. Relational memory theory (Cohen & Eichenbaum, 1993), on the other hand, very directly addresses how memory for arbitrary relations might form a flexible representation (in the graph-like manner discussed previously). However, even once we accept that a graph-like representation is at the core of memory for reconstruction, the problem is still far from solved. We don't know, for instance, if a subset of items is being represented relationally, bound to each other, while a single item is being bound to some aspect of the environment in exclusion of the other items. Additionally, the presence of higher-order organizational information (such as contextual information) might bias the reconstruction in ways which are measurable in the change of position of the groups of entities which share that higher-order property. The exact definition of context, however, is a topic of heavy debate and will not be addressed in this work.

Now, for identity information, we also have several possibilities. If the identity information is being convolved with some aspect of the spatial information, we would expect it to be subject to the same transformation and invariance rules outlined in the previous paragraph. If, however, the items are being bound to particular locations (as was suggested in Watson et al., 2013), these bindings might be subject to their own invariance rules. This situation is more likely if the items are arbitrary, i.e. they just as easily could have been any item. Arbitrary items cannot be as easily convolved with location, as no expectancy can be formed about what item should be associated with a given location (thus, arbitrary). In reality, arbitrary items may still end up with alternative associations by chance (for instance, if an image of a plane happens to appear over an image of a tree, this obviously has a location-identity association which is non-arbitrary as planes tend to fly over trees, but it may have occurred by chance), but if the items are arbitrary, we can imagine them, in some way, as a disjoint domain of their own which must be bound to the continuous, connected domain of space. This binding may not be damaged by transformation of the

space as it is being done via creating arbitrary relations from single points to identities. The various

identities, however, could end up confused in some way if that information was not remembered but

the location information was. In this way, particular identities might be assigned to the wrong location,

or identities might switch places (see **Figure 1.6** for a visual description of these representational

differences). If the relations are arbitrary, these errors should be thought of as distinct from the sorts

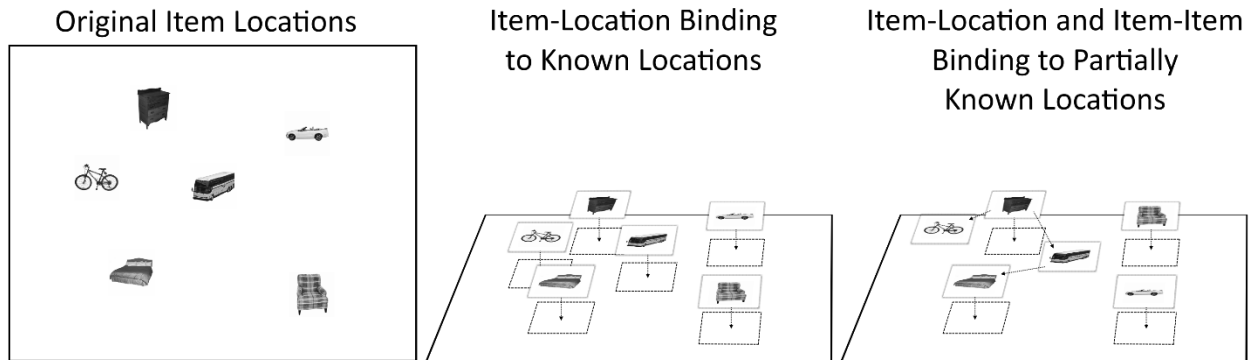which might result from some failure to remember location.

| Original Item Locations | Item-Location Binding to Known Locations | Item-Location and Item-Item Binding to Partially Known Locations |

Figure 1.6: Examples of different identity binding methods. Item-Location Bindings to Known Locations assumes that item identities are being associated to particular locations and the locations are being remembered independently. Item-Location and Item-Item Binding to Partially Known Locations shows how, even in the absence of perfect location information, item-item relations can allow imprecise reconstruction of locations. Similarly, item-location binding can still be present, but identity can be misassigned even with perfect memory for location (i.e. the chair and the car are switched). Although we cannot easily know the particular representation being used by an individual on a particular trial, we can observe when these various errors occur and narrow down the possible representations and the information which is lacking.

1.1.4 Sampling and Encoding During Study

Finally, the discussion up until now has been exclusive to the representational portion of memory, but it

largely neglects the learning aspect of reconstruction. When performing a reconstruction task,

individuals are given some amount of time to study the information. During this time, there are

numerous behaviors which might be performed in order to attempt to gather the information. The

structure of human sensory systems is such that we generally sample specific aspects of information

one-at-a-time, rather than assimilating the whole environment simultaneously (see Kveraga & Bar, 2015

for a thorough treatment of the topic). This sampling behavior is of particular interest as it is often

measurable. If the information being studied is all visual (such as a computer screen task), the eyes will

be the only sampling medium and, as such, the movement of the eyes in some way determines the

information which can be encoded. Eye movements have been successfully used to identify critical

differences between viewing patterns in hippocampal damaged individuals and comparison participants due to memory (Cohen & Banich, 2003; Hannula & Ranganath, 2009). In Virtual Reality or Real-World environments, navigation (the combination of body-movement and eye-movement) provides a more flexible form of sampling (and, in fact, eye movements can be a subset of this sampling as the technology to sample eye movement at the same time as head movement and environment position has been available for nearly 20 years; Duchowski et al., 2000). In either case, sampling comes down to a stream of information which is received by the system whose contents can be guided by in-the-moment demands; in fact the Covert, Rapid, Action-Memory Simulation (CRAMS) model specifically addresses this sort of rapid, in-the-moment, memory-guided decision making, suggesting that prefrontal-hippocampal interactions provide rapid simulations of potential outcomes of decisions in order to guide behavior (Wang et al., 2015). If the task at hand has few demands, very little sampling will be required, however, in a demanding task, extensive sampling may be required to acquire the necessary information. Tasks are sometimes designed to require participants to sample particular aspects of an environment (via masking all but some small region of the environment, *a la* Yee, 2012, or by forcing the participant to acknowledge sampling of particular aspects of the environment before proceeding, for example), and these demands may also bias behavior in particular ways. Sampling is a complicated topic, and as such, it will only be examined in relation to the test performance and not in a more holistic, information rich manner.

1.2 Overview of Chapters

The remainder of this thesis will discuss algorithmic and mathematical formulations of domains and entities in reconstruction tasks, starting specifically with the aforementioned example of 2D space, validating the methods in hippocampal damaged patients. These methods will be used to show that hippocampal damage does uniquely impair some, but not all, measures of performance in reconstruction (Chapter 2). Later sections will show how these methods can be applied to a temporal domain and show how, despite its differences, temporal memory representations may be treated in much the same way as spatial memory representations (Chapter 3) under this theoretical framework. Temporal and spatial domains may share many similarities, but errors in each may be distinct in other ways which can be seen in reconstruction performance. In particular, contextual information may have a similar influence on memory in both space and time, when it comes to precise remembering of location information, but relational information may be in some way different in these two domains. Finally, we will examine aspects of sampling of a domain which might influence the ability to remember and reconstruct information (Chapter 4), and we will specifically demonstrate how changes in navigation and

exploration behavior relate to changes in various memory measures such that rapid changes in systematicity and complexity of navigation relate to faster improvement of subsequent relational memory performance, and faster increases in temporal-contiguity of navigation relates to faster improvements in several aspects of relational and contextual memory.

1.2.1 Reconstructing Relational Information

In the second Chapter, a computational framework based on the theories outlined so far in the introduction is established which attempts to approach the task of inferring information about spatial reconstruction performance in hippocampal damaged patients from a principled perspective, making as few assumptions as possible and carefully defining different *types* of spatial relational information which are quantifiable in a reconstruction. Previous work has identified misplacement (i.e. the sum of Euclidean distance of the reconstructed locations of items to their target locations) as an error which shows greater magnitude in hippocampal damaged patients than comparison participants (Huttenlocher & Presson, 1979), and more recent work has hypothesized that some of this may be due to swapping of item identities between otherwise remembered locations (Watson et al., 2013). The work here expands on these ideas with a far more precise mathematical formulation of errors involving misassignments of items, global transformations (i.e. translation, scaling, and rotation), and swaps/cycles of items which are misassigned as a group. Three hippocampal damaged patients are evaluated in a computer-based spatial reconstruction task with 6 unique items and 32 trials. Nine age, sex, and education matched comparison participants perform the same task. In agreement with previous theories of relational memory and hippocampal damage (Cohen & Eichenbaum, 1993; Eichenbaum & Cohen, 2001; Konkel & Cohen, 2009; Konkel et al., 2008), single-item misassignments (i.e. placing of an item in another item's location without any presumption about where the item which should've occupied that location was placed) are significantly more common in hippocampal damaged individuals than comparisons. However, patients and comparisons do not show a difference in the number of overall locations which are correctly reconstructed with some item (though not necessarily the correct item identity). This suggests that hippocampal patients are able to maintain certain relational information (i.e. general studied locations), potentially via an alternate representation that leverages other brain regions and/or memory systems.

Furthermore, despite previously identified differences in swapping between patient and comparison participants, with these more sensitive measures, we show that this difference is only present at low set-sizes (i.e. few items, <=4, in the reconstruction). At 5 items, comparison participants begin to commit

identity-location spatial relational errors and, at this point, we find that no difference is present between patients and comparisons in the number of compound errors (swaps and cycles of items). Additionally, although identity information does account for the majority of the difference between patient and comparison participants, a significant difference remains in their misplacement even after all corrections are made. This "local" misplacement may represent a deficit in precision as described by others (Kolarik et al., 2016, 2017; Yonelinas, 2013).

This work serves two primary purposes: first, the paper establishes a methodological framework for making inferences about spatial reconstruction performance (and, in general, reconstruction performance in any domain) which makes minimal assumptions and approaches the issue from a principled perspective. Second, the work clarifies the previous theoretical accounts of relational memory and the hippocampus by showing that it is the arbitrary identity-location information (i.e. the information which "could have just as easily been anything") which is impaired in hippocampal damaged patients, not the overall location information. This work speculates as to how such information might be maintained or represented, but future studies which are not part of this work will be necessary to clarify how some relational information might still be maintained via alternate mechanisms.

1.2.2 Memory during Time Travel: Spatiotemporal Navigation, Contextual Boundaries, and Relational Memory Errors in Virtual Reality

Beyond just the spatial domain, episodic memory (which is critically supported by the hippocampus; Tulving, 2002a) involves the temporal domain as well. There is an extensive literature studying the temporal domain showing that humans naturally segment events into discrete regions in time (Zacks et al., 2007), that contextual boundaries impact perceived distances across time (DuBrow & Davachi, 2013), that order memory is impacted by recency, contiguity, and primacy effects (Howard et al., 2015; Murdock, 1962), and an enormous litany of other discoveries, all which make the same crucial assumption – that time need always be sampled unidirectionally and in isolation of space, despite the fact that space and time are being constantly sampled together and memory for temporal events may involve more flexible representations than just a unidirectional sampling allows us to examine. The work presented in this section shows that this assumption is not necessary and space and time can be put on a more equal footing, allowing free navigation of both domains simultaneously (known as a Spatiotemporal Navigation Task) in Virtual Reality (VR). Virtual Reality allows increased measurement and control capabilities over other methods, increasing a sense of presence, and providing greater ecological validity than more traditional methods (Kuliga et al., 2015; Schultheis et al., 2002). Many of the same effects seen in more traditional tasks can be seen in this Spatiotemporal Navigation Task

(context boundary effects – in which items which share context are remembered as closer together and items across contexts further apart; misassignments – where one event (i.e. an entity whose identity is associated with a specific spatiotemporal location) is placed in another event's location; overall misplacement improvement (i.e. the sum of the Euclidean distances of all events to their targets) due to repeated study of the same events. While many of these effects are similar to previous tasks, despite the present study allowing for the ability to traverse the timeline freely, some unique differences can be found between space and time when comparing relational information in each domain.

Leveraging the methodological work from Chapter 2, Chapter 3 shows that misassignments in time (i.e. placing an event in another event's temporal location) is far more likely than misassignment in space (i.e. placing an event in another event's spatial location) in healthy young adults even with the ability to freely explore both domains. This difference is persistent across trials, with spatial misassignment dropping to near perfect accuracy by the last trial while temporal misassignment remains significantly higher. However, a significant linear trend is present in the temporal misassignment across trial suggesting it might eventually plateau if more trials were utilized. The ability to improve an understanding of the relational information contained within the temporal component of the set of events is particularly interesting in light of work by Howard Eichenbaum and colleagues showing the existence of time cells in the hippocampus which fire in relation to specific points in time (Eichenbaum, 2014; Kraus et al., 2013). It is unclear if these time cells are firing in accordance with any temporally task-relevant stimuli, relating to a topological representation of time (i.e. where the relative distances between events are represented independently of a grid-like representation), or firing exclusively in relation to some task start point (as is the experiments up until now), relating to a metric (i.e. grid-like) representation of time (see Ekstrom & Ranganath, 2017 for a discussion of the distinction between topological and metric representations). Based on what we know of other selective cells in the hippocampus (i.e. orientation selective), we would predict that time cells should exist which fire relative to task-relevant temporal stimuli, i.e. a topological representation of time. However, without the manipulation allowing for the re-exploration of temporal moments, it is difficult to dissociate the two possibilities as temporal relations always occur at precisely the same temporal distance from one another. This task presents a manipulation which might one day lead to a better understanding of these temporal representations, though this behavioral evidence, alone, will not be sufficient to disentangle these possibilities.

This task does not only address spatial and temporal reconstruction; it also contains temporal contextual information in the form of changing background colors in the environment at different moments in time. Using these temporal contexts, temporal context boundary effects can be assessed. Although these effects have been looked at for temporal order memory in previous experiments (DuBrow & Davachi, 2013; Ezzyat & Davachi, 2014), those experiments never allowed the participant to freely transition between contexts at will, to (re)explore contexts as much as they would like, or retested on the same information several times (to assess improvements to contextual vs. overall memory). In this task, participants repeated the same event sequence 4 times, and showed improvements to their overall memory on each subsequent trial. However, despite improvements in performance across nearly every other metric, temporal context boundary effects did not diminish across trials. This strongly suggests that this "error" in placement is, in fact, a bias created by an organizational principle which uses context to separate groups of events from one another. This suggestion is consistent with other computational models of the impact of temporal context on memory (Howard et al., 2014); however, these models do not generally allow for the assumption of free traversal of time.

Interestingly, this contextual bias does not just impact memory for the relative distance between events, but contextual information also biases relational information in time (i.e. temporal misassignments, from earlier in this section). Across all trials, a significantly higher number of misassignments are to same-context locations that should occur by chance. Moreover, this effect actually increases across trials as the impact of context on distance judgements holds steady and the overall number of misassignments decreases, suggesting that this bias is not easily corrected with restudy and reflects an organizational principle which causes within-context items to be associated to each other but not necessarily with the same precision on an individual-item level.

Together, this task and the mathematical framework from Chapter 2 serve as a critical first step in forming a more comprehensive understanding of episodic memory and relational memory in space and time, simultaneously, showing that, when we put space and time on a more equal footing (by allowing participants to freely sample both domains), we can study memory for events with more flexibility and fidelity to illuminate critical elements of memory organization. In particular, we show that relational memory in space and time may not be represented with precisely the same organization (as shown by drastically different relational memory errors in each domain) and that contextual information creates systematic biases in both distance judgements and relational memory errors which do not necessarily improve with restudy.

1.2.3 Spatiotemporal Navigation, Sampling, and Information Encoding in Virtual Reality

Previous sections have primarily highlighted test results, while leaving study-time behavior largely unaddressed. Chapter 4 examines various measures of study-time performance in the most complex case of the tasks discussed in this work, the Spatiotemporal Navigation Task, and shows that changes in measures of complexity, systematicity, and contiguity of navigation relate to changes in relational and contextual memory performance. Previous work on eye tracking in hippocampally dependent tasks and hippocampal damaged populations has revealed a variety of visual sampling phenomena which are predictive of aspects of hippocampally dependent memory (Hannula, 2010; Hannula & Ranganath, 2009; Hannula et al., 2007). However, sampling in a 2D environment via eye movement is somewhat different from a computational perspective than navigation within a 4D environment. Basic navigation metrics such as the distance travelled in space and time have been evaluated, showing that performance on these metrics improves across trials in the Spatiotemporal Navigation Task and the rate of that improvement relates to the rate of improvement in the misplacement in each domain (i.e. faster reduction in distance relates to faster reduction in misplacement). Additional metrics of complexity of navigation from other work in spatial navigation (such as the Fractal Dimension measure; Daugherty et al., 2015) and novel metrics which relate to systematicity of navigation (i.e. Lacunarity) are specifically examined to observe relationships with relational memory and contextual biases at test. Additionally, an analysis of the order of navigation versus order of reconstruction and associated accuracies (more akin to the Recency and Contiguity effects explored in traditional temporal free recall tasks) is examined.

In this task, individual improvements in spatial and temporal navigation are shown to relate to improvements in memory in those domains separably, suggesting that spatial and temporal representations may, in some way, be separable in this task; relational memory improvements are shown to be uniquely tied to changes in navigation complexity and systematicity, suggesting a critical and complex interplay between in-the-moment, memory-guided decision making and subsequent relational memory efficacy; context boundaries are suggested to act as more of a discriminatory feature (at least in this task) than one used to strengthen within-context relational memory organization accuracy as there is a significant relationship between changes in context boundary crossing and both the context boundary effect and across-context relational memory errors; and a preference towards exploring an otherwise temporally-flexible environment in the implied, forward order with increasing contiguity is suggested to be a critical element in improving temporal, relational, and contextual memory organization.

Chapter 2: Reconstructing Relational Information[1]

2.1 Introduction

Relational memory is critically supported by the hippocampus (Cohen & Eichenbaum, 1993; Eichenbaum & Cohen, 2001; Konkel & Cohen, 2009; Konkel et al., 2008). For example, memory for relationships of items to spatial locations is highly sensitive to hippocampal damage in humans (Hartley et al., 2007; Allen et al., 2014; Smith & Milner, 1981a; Watson et al., 2013). Moreover, it has been shown previously that the hippocampus is necessary for all "manner" of relational memory (Konkel et al., 2008); that is, it is critical for binding information regardless of the informational domain, be it spatial, temporal, sequential, associative, etc. However, while the hippocampus can be involved regardless of the domain of information, there are many possible relations between the constituent elements within a given domain, such as item identities (i.e. labels that distinguish items from one another, such as names), locations (i.e. positions in the domain in which an item was present during study), and environmental elements (i.e. the boundary of the environment or landmarks). It is possible that not all of the relational information among these elements critically relies on the hippocampus, and that some *types* of relational information may not be impaired by hippocampal damage.

Importantly, we distinguish between a *manner* of relation and *type* of relation. Hippocampal involvement in all *manner* of relation means, more generally, that the hippocampus can be involved regardless of the informational domains (i.e. spatial, temporal, social, associative, etc.). However, within a given domain, particular *types* of relational information exist (and are enumerated here) which may not all be equivalent in terms of hippocampal involvement. It is possible that when comparing *types* of relations in two different domains, seemingly similar relations may have distinct representations with different amounts of reliance on specific memory systems. For instance, item-item relations within a spatial domain containing several items can be used to derive information about item locations even with an imperfect representation of each individual relation (this point will be elaborated upon later), while item-item relations in an arbitrary word pair cannot be easily used to derive additional information about other arbitrary word pairs. Similarly, item-environment relations in some spatial-relational tasks may or may not be informationally equivalent to item-environment relations in others (such as studying the location of a single item in a scene vs. the locations of many items within an empty

---

environment). A primary reason for this difficulty is in the diversity of uses of terms like 'item' and 'environment' across different tasks.

It is the aim of this work to precisely define *types* of relations within the spatial domain in such a way as to determine if memory for specific *types* of spatial relations are impaired (or not) by damage to the hippocampus. To this end, we propose a framework for systematically classifying *types* of relations in a spatial memory task. This framework allows us to: 1) distinguish multiple types of first-order (i.e. pairwise or one-to-one) and higher-order (i.e. group-wise or many-to-many) spatial relations, and 2) determine if these various types of spatial relations are differentially impaired by hippocampal damage by inferring the presence or absence of relational information via observations of different types of errors in a reconstruction.

Here we take advantage of the rich data generated by spatial reconstruction (SR) paradigms, in which multiple items are studied in various spatial locations before participants are asked to reconstruct (i.e. freely place) each item in its remembered location. There is a long tradition of using SR paradigms to study spatial memory (Huttenlocher & Presson, 1979) and its susceptibility to hippocampal damage (Jeneson et al., 2010; Smith & Milner, 1981; Watson et al., 2013). These experiments have historically used a general quantification of "misplacement error" in space by calculating the sum of the Euclidean distance between each placed item and its studied location. The typical finding of these experiments is that individuals with hippocampal damage show increased overall misplacement relative to comparison participants.

One critique of the use of misplacement as the sole metric of performance on SR tasks has been that it provides relatively little information about the nature of spatial memory deficits following hippocampal damage (Watson et al, 2013). In other spatial paradigms, such as the virtual Morris Water Maze (vMWM), some theories of hippocampal function such as the Precision and Binding Model (PBM; Yonelinas, 2013) would predict the increased misplacement seen in SR tasks, with PBM, in particular, predicting poorer precision of spatial memory in hippocampal damaged patients (Kolarik et al., 2016, 2017). This could be consistent with the explanation that high levels of misplacement could theoretically reflect an inability to remember coordinates on a "grid-like" mental representation of the display, with patients showing a deficit in the resolution of the grid. Alternatively, high misplacement could reflect a deficit in representing inter-object configural or relational information (among other possibilities, discussed below). Due to this resolution issue in using misplacement as a primary metric of performance, there have been recent attempts to develop additional metrics of spatial reconstruction

performance that help to disentangle these possibilities (Jeneson et al., 2010; Watson et al., 2013). In particular, Watson et al., (2013) focused on a very specific type of spatial relational error termed a "swap" error. "Swap" errors are errors in which the relative position on each x and y axis for a pair of items is flipped. Such errors appear to be strongly influenced by hippocampal damage. In fact, patients with hippocampal lesions even make "swap" errors in set sizes as small as two items and on surprisingly short time scales (Watson et al., 2013).

These data suggest that when the hippocampus is impaired, relational binding errors occur regardless of the number of items to be remembered (e.g., in a set size invariant manner). Moreover, swap errors in Watson et al (2013) were found to be a significant contributor to the overall differences in misplacement rates between patients and healthy comparison participants. However, there are many additional potential relation-based explanations for increased misplacement in hippocampal patients. Patients could be making global errors of some sort (i.e. moving all items down, squishing items together, or rotating all the items), local or "noisy" errors (i.e. placing items pseudo-randomly based on a heuristic or limited memory for locations), or relational errors (i.e. placing items in the wrong location given their identity, but that location would have been valid for a different item). Thus, in the present study, rather than focus on one type of spatial relational error (e.g. swaps), we sought to identify the multiple types of relations in space, to organize them in a systematic framework, and to determine if all, or only some, types of relations are impacted by hippocampal damage.

At the core of our proposed framework is the classification of three primary types of first-order spatial relations: item-environment relations (e.g., one item is in the upper left corner of the display), item-item relations (e.g., one item is below and to the left of another item), and identity-location relations (e.g., item A belongs in location A). Note that terminologically, the use of the word "item" here is specific to location information (with identity information being separated out as identity-location information). An analogy to help clarify the meaning of these different relations can be found in navigating a shopping mall. While navigating, it may be important to remember the locations of stores. This would require you to remember location information about where the stores of interest are in relation to the shopping mall layout (i.e. item-environment relations) or perhaps the locations relative to other stores (i.e. item-item relations) as well as which store is in which particular location (i.e. identity-location relations). The hippocampus has been implicated for decades in item-environment relations. Indeed, animal models demonstrate hippocampal place cell firing tied to the distance and direction of item cues in an environment (Gothard et al., 1996; John O'Keefe et al., 1996). While humans with hippocampal damage

are able to retain some item-location information via maintenance in working memory, this maintenance is transient (Allen et al., 2014; Libby et al., 2014). In both humans and animals, arbitrary item-item association has been extensively studied. Memory for item-item relations is impaired following selective damage to the medial temporal lobe in rodents (Bunsey & Eichenbaum, 1993), monkeys (Murray et al., 1993), and humans (Giovanello et al., 2003). Furthermore, humans with hippocampal damage show item-item impairments for all manner of relations (e.g. spatial, temporal, associative, etc.; Konkel et al., 2008).

In addition to these first order relations, we examine higher-order, compound relations (e.g., item A, B, and C form a group), which can contribute to compound spatial relational memory errors (e.g., the group is translated to the left, all of the items are squished towards the group's center, or two items swap locations). Compound relational errors can be seen in two forms in our analysis. Firstly, swaps of items (i.e., when two items are placed in each other's location) and cycles of items (i.e. when more than two items are placed in each other's locations) are forms of compound error primarily involving multiple identity-location errors. Secondly, various transformation errors (translation, scaling, and rotation) are forms of compound errors involving multiple item-item and/or item-environment errors.

In addition to first-order (i.e. pairwise) and higher-order (i.e. compound) relations, we can begin to consider the question of how general, Gestalten shape information, i.e., configural features which constitute parts of a unified whole, may be used independently of the hippocampus to maintain more global spatial information. While the current data cannot provide a definitive response to this proposal, the relevant data are suggestive and considered.

Finally, given the evidence that hippocampal damage impairs reconstruction performance even at very small set sizes (Watson et al., 2013), the implication of varying set size on memory for these different types of spatial relational information was considered. To this end, we take advantage of an existing dataset (Watson et al., 2013) that recorded SR task performance from hippocampal patients across set sizes (i.e., from 2 items to 5 items) and reanalyze those data using the current analysis framework. We first evaluate the extent to which hippocampal patients and comparison participants committed relational memory errors of particular types even at small set sizes (e.g., 2 items). Second, we determine whether those error types change or remain consistent as set size increased.

In summary, in the work presented here, we investigate the types of relations that are sensitive to hippocampal damage from a first principles approach by breaking down reconstruction errors in an SR

task into multiple error types, which can then be evaluated simultaneously. We consider three types of first-order relations (i.e., item-environment, item-item, and identity-location relations) as well as two types of higher-order compound relations (i.e., swaps and cycles). We evaluate the degree to which the overall difference in misplacement between hippocampal patients and matched comparisons can be explained by these distinct error types. Finally, we discuss the implications of these findings for theories of hippocampal function.

2.2 Materials and Methods

2.2.1 Participants

Behavioral data were collected from 3 patients with hippocampal damaged and profound declarative memory impairment or amnesia (see **Table 2.1** for details) and 9 healthy comparison participants. Three of the healthy comparison participants were matched to each hippocampal patient for sex, age (+/- 5 years), and education (+/- 2 years).

DEMOGRAPHICS						NEUROANATOMY		NEUROPSYCHOLOGICAL						
SUB.	Age	Edu.	Sex	Hand.	Etiology	Hipp.	Lesion	WAIS-III FSIQ	WMS-III GMI	WMS-III WMI	BNT	TT	CFT Copy	WCS Cat
1846	52	14	F	R	Anoxia	-4.23	Bilateral HC	84	57	85	43	41	28	6
2563	60	16	M	L	Anoxia	N/A	Bilateral HC	102	75	99	52	44	36	6
1951	63	16	M	R	HSE	-8.10	Bilateral HC + MTL	106	57	108	49	44	32	6
MEAN	58	15	-	-	-	-	-	97	63	97	48	43	32	6

Table 2.1: Demographic, neuroanatomical, and neuropsychological characteristics of participants with hippocampal amnesia. Hand. = Handedness. Edu. = years of education. HSE = Herpes Simplex Encephalitis. HC = hippocampus. + MTL = damage extending into the greater medial temporal lobes. HC Volume = hippocampal volumetric z-scores as measured through high resolution volumetric MRI and compared to a matched healthy comparison group (Allen et al., 2006; Buchanan, 2005). WAIS-III FSIQ = Wechsler Adult Intelligence Scale–III Full Scale Intelligence Quotient. WMS-III GMI = Wechsler Memory Scale–III General Memory Index. BNT = Boston Naming Test. TT = Token Test. CFT = Complex Figure Test; WCT = Wisconsin Card Sorting Task; Cat = Number of categories achieved out of six.

Two patients experienced anoxic/hypoxic episodes (1846, 2563) resulting in bilateral hippocampal damage and the third patient contracted herpes simplex encephalitis (1951) leading to more extensive bilateral MTL damage affecting the hippocampus, amygdala, and surrounding cortices (**Figure 2.1**). Structural MRI examinations completed on 2 of the 3 patients confirmed bilateral hippocampal damage and volumetric analyses revealed significantly reduced hippocampal volumes. Participant 2563 wears a pacemaker and was unable to undergo MRI examination and thus their damage was confirmed by computerized tomography; damage was confirmed to the hippocampus.
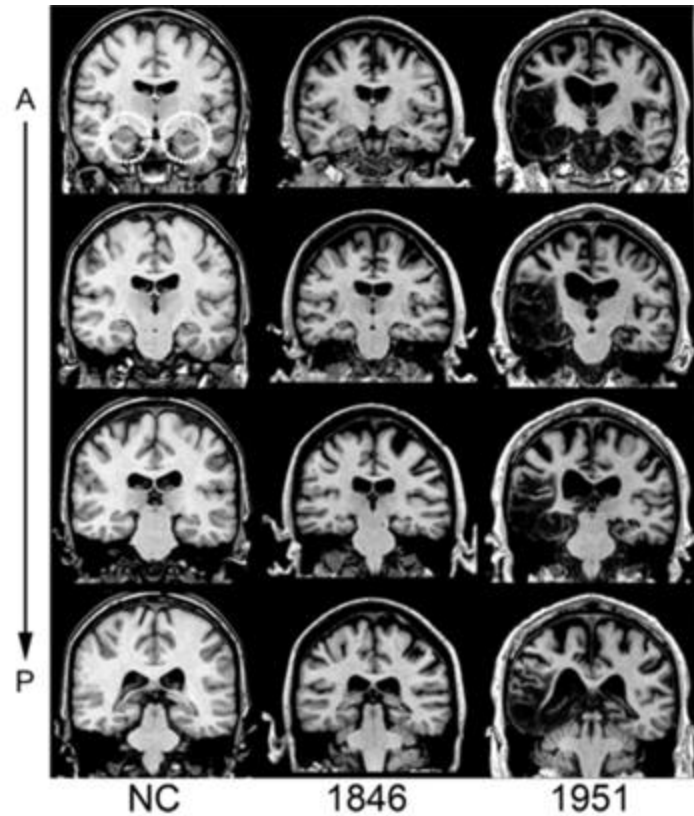
Figure 2.1: Magnetic resonance scans of hippocampal patients. Images are coronal slices through four points along the hippocampus from T1-weighed scans. Volume changes can be noted in the hippocampal region for patient 1846, and significant bilateral MTL damage including the hippocampus can be noted in patient 1951. A = anterior, P = posterior, NC = healthy comparison brain."

Neuropsychological examination confirmed severe declarative memory impairment in all hippocampal patients (M =63; Wechsler Memory Scale-III General Memory Index, more than 2 SDs below population norms) in the context of within normal performance of standardized measures of intelligence, language, visual perception, working memory, and executive functioning (see **Table 2.1**).

2.2.2 Experimental Paradigm

Participants completed a computerized spatial reconstruction task (Monti et al., 2015; see top of **Figure 2.2**). Participants studied 6 items (gray scale, nameable objects; Brodeur et al., 2014) per trial, arranged pseudo-randomly in an 950×600 pixel area within a 1280×1024 pixel computer screen across 32 total trials. During the study phase, participants viewed the 6 items for 16 seconds, followed by a 5 second delay (blank screen). At test, which began immediately after the delay, items were randomly placed in a line at the top of the screen, and participants used a mouse to drag each item to its remembered location on the screen. Participants had unlimited time to reconstruct the studied display. When they were finished, participants pressed space bar to start the next study trial. Participants' eye

25

movements were recorded throughout the experiment. Eye tracking data are not reported here and will be reported elsewhere (Lucas, Duff, & Cohen, in preparation).
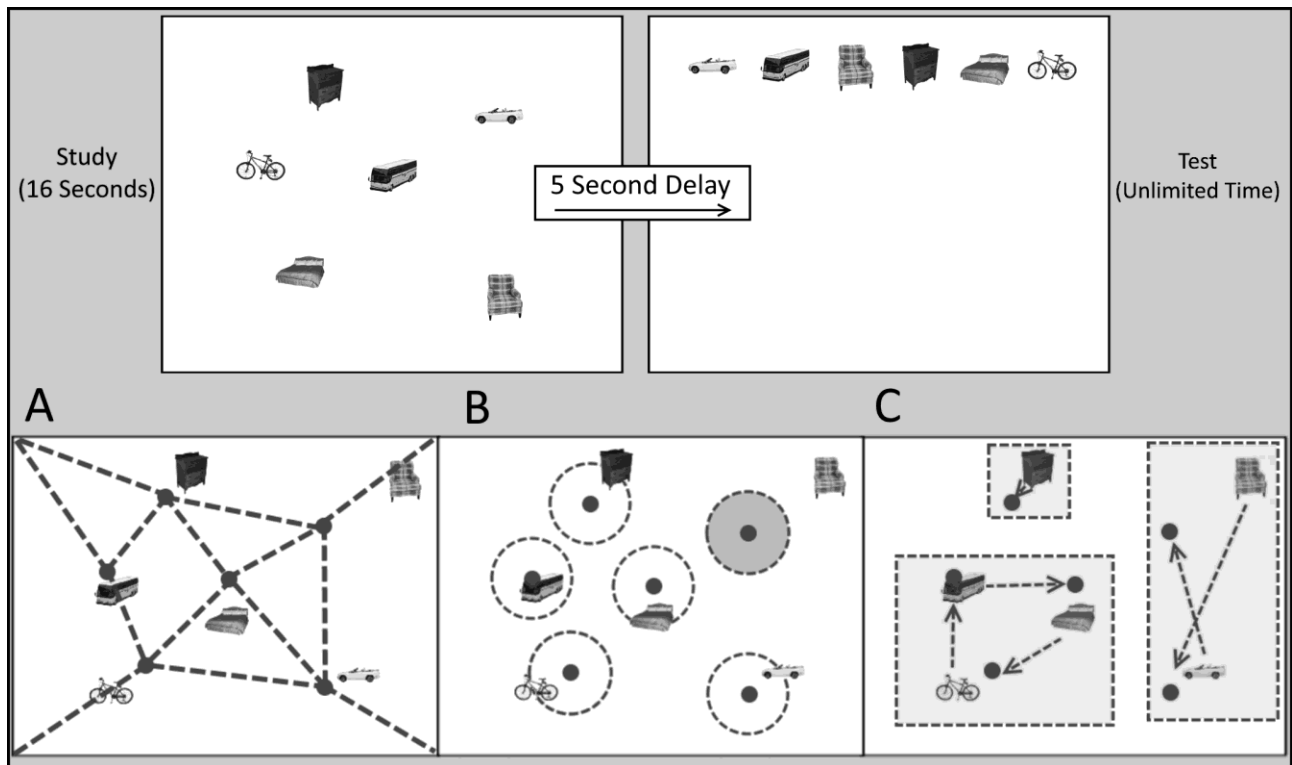

Figure 2.2: A sample trial from the spatial reconstruction task; A, B, and C show different interpretations of the same reconstruction where the objects are in the reconstructed locations and the dots are the target locations from the study phase. (A) shows item-item and item-environment relations (dotted lines), which might be used to determine the locations of the dots. (B) shows reconstructed item to target location relations where gray shaded circles indicate a failure in placement of items to those locations (note that identity is ignored as, for example, the bus is not in its target location). (C) shows compound relations where some items roughly swap locations (although not necessarily precisely, in this case being a "partial swap"), some show cycle relationships (i.e. Bus→Bed→Bicycle→Bus), and some show single-item placement in the correct target location.

2.2.3 Analysis Methods

To determine if item-environment, item-item, identity-location, and compound relations showed selective differences in performance between patients and comparison participants, a new set of tools for analyzing spatial reconstruction data were created. These methodologies were designed with the intention of making as few assumptions as possible about the nature of the reconstruction to separate the types of errors via more mathematically rigorous and interpretable metrics. This approach applies Point Set Registration methods (see Besi & Mckay, 1992 for an introduction to the topic), in which a set of points (often derived from key points in an image) are assigned to an independent set of points (often from a reference image) based on an inferred relationship between the two sets. The SR paradigm is analogous to the Point Set Registration process of understanding the relationship between

reconstructed points and their target (i.e., studied) counterparts, with two primary differences. Firstly, in the case of Point Set Registration, it is often not guaranteed that every point in the data set will correspond one-to-one with a point in the target. In the case of the SR task, we have a much stronger reason to assume that such a one-to-one correspondence should exist (though it may be interesting to consider the sorts of memory errors that might cause this assumption to be violated). Secondly, in Point Set Registration, it is often the case that there are many more points from which to define a model than exist in the SR task. For this second reason, especially, it is important that we proceed with caution in what models we apply to our data to avoid overfitting to a small set of points.

The model for the comparison of the reconstructed and target (i.e., studied) points can be thought of as having four primary components: identity remapping (wherein the identity of each item is removed and the set of placed item locations is mapped onto the target locations in a way which minimizes error), global error correction (wherein global translation, scaling, and rotation errors are subtracted from the placed item locations), location placement evaluation (wherein a binary determination of accuracy is made for each location), and compound error evaluation (wherein "swaps," in which two items are misassigned to each other's locations, and "cycles," in which more than two items are misassigned to each other's locations, are identified).

2.2.3.1 Identity Remapping

For this first step, identity remapping, a one-to-one mapping between the reconstructed points and targets must be created. This mapping can be determined by a variety of methods, each with their own advantages and disadvantages (Burkard et al., 2012); however, for the sake of simplicity and interpretability, it is often preferable to find the mapping which globally minimizes the overall error in the reconstruction (effectively eliminating any effects of item identity-location misassociation in the reconstruction). This can be solved using the following Assignment Problem (Kuhn, 1956; Kuhn, 2010; Munkres, 1957) equation:

$$\min \sum_i \sum_j C_{ij} X_{ij}$$

where C_{ij} is the cost matrix associating each vertex in a graph and X_{ij} is the binary assignment of an undirected edge to the graph. Using the above equation, we find the set of X_{ij} associations (which we will call "the map") which results in the smallest numerical value, forming a minimal one-to-one mapping of the placed item location to a target (studied) location.

2.2.3.2 Global Error Transformations

Once a one-to-one mapping has been found, the next primary component to be determined is the global error transformation which finds shared error across all the items. This transform could be any function, and with enough parameters, the items could be mapped perfectly onto their assigned target locations. However, because our intention is to differentiate between global relational errors and item-level errors, a simplified transformation is used such that only translation, scaling, and rotational errors are included in the model. Together, these transformation components describe a 4 parameter system (x translation, y translation, scale, and rotation), which can be used to model global error across all items in a trial (cf. Similarity Transform; Cederberg, 2001). One could speculate as to what might cause these different types of errors (i.e. translation being a misrepresentation of the entire stimulus space in the environment with an offset, scaling being a squishing or stretching of the stimulus space, and rotation being a misrepresentation of the directions in the environment). Regardless, by identifying this global error, we can subtract out its effects from each item's location so as to get a more precise measure of individual item accuracy independent of shared, systematic, spatial errors. The similarity transform can be more precisely stated as a solution to the equation:

$$f(x) = rAx + t$$

where r is the ratio of similarity (or scaling factor), A is an $n \times n$ orthogonal matrix and t is a translation vector. A can be decomposed into a single θ value which represents the magnitude of rotation around the point set center. This transformation can be computed using Umeyama's algorithm (Umeyama, 1991) which uses Singular Value Decomposition (Cline & Dhillon, 2006; Golub & Kahan, 1965) to obtain a transformation matrix. It is useful to note that it is not always possible to find a transformation given a set of points if they do not actually share any global error of the types we've specified (i.e. translation, scaling, and rotation). In our model, we allow rotation to fail independently of translation and scaling, as it is more difficult to find rotational components in small data sets. If all attempts to find a transformation fail, no global adjustments are made and the transformation data from those trials are excluded from the analysis. Singular Value Decomposition computes the transformation matrix by solving the following equation:

$$M = U\sum V^*$$

where M is the dot product of the reconstructed and target point locations divided by the number of points, U and V^* are unitary matrices, and $\sum$ is a diagonal matrix of non-negative, real numbers.

Umeyama's algorithm allows the decomposition of the matrix $\boldsymbol{M}$ into the translation, scaling, and rotation components.

2.2.3.3 Location Placement Evaluation

Now that associative and global errors have been subtracted from the point locations, a measure of reconstructed-to-target location placement accuracy can be evaluated in a manner that is not influenced by incorrect item-identity choices or global rearrangement. Previous studies have attempted to include accurate placement as an aspect of the analysis using an arbitrarily chosen distance value (Jeneson et al., 2010). However, an inflexible threshold imposes a very strict, nonlinear boundary on our evaluation of accuracy, which is difficult to handle without training participants to operate within this constraint (potentially influencing behavior in unintentional ways). It would be preferable for such a boundary decision to be made flexibly, based on the variability in local misplacement (i.e., misplacement that is not shared amongst the items nor due to item identity errors). To accomplish this goal, we determine accuracy based on the data by creating a statistical distribution of the misplacement after correcting for item-identity and global error and use the confidence intervals for the distribution (in this case, 95% confidence intervals) as the accuracy threshold. This distribution could be formed on a trial-by-trial basis, but it is more stable to compute it across all trials for an individual as there are more sample points (192 points instead of 6). This trade-off means that the threshold will be slightly stricter than the data would suggest for earlier trials and slightly less strict for later trials due to the improvements in accuracy normally associated with learning the studied locations. The confidence interval equation:

$$d < \bar{x} + z^* \frac{\sigma}{\sqrt{n}}$$

is used to determine if an item is within the accuracy threshold, where d is the distance of a given item to its target, $\bar{x}$ is the mean distance of all items to their targets, σ is the standard deviation of all items to their targets, n is the number of items, and z^* is a chosen z-score (usually 1.96 for 95% confidence intervals). This choice trades a rigid, predetermined threshold condition for a new assumption: that the participant placement of items will generally be in the studied locations (to a degree specified by the confidence interval width). This assumption may not be correct for patients with hippocampal damage as overall misplacement for patients is greater than for matched comparisons (Huttenlocher & Presson, 1979; Jeneson et al., 2010; Smith & Milner, 1981a). Therefore, we opt to use the accuracy thresholds computed from the comparisons as the measure of accuracy for the patients. This choice allows our definition of placement accuracy to be unified between the groups, with the matched comparisons

29

determining their associated patient's placement accuracy threshold. This adjustment is necessary in deciding a threshold for patients because we are forming a statistical distribution based on individual item misplacement which incorporates all the individual item variability. This item variability is not necessarily matched between patients and comparisons given that the comparisons generally perform better. Thus, we use the stricter criteria, the comparison thresholds, for placement accuracy in patients.

2.2.3.4 Compound Identity Error Evaluation

Finally, once all item-location assignment and global errors have been subtracted from the data and placement location accuracies have been determined, item-to-location mapping can be examined for both accurate locations and inaccurate locations. Using the mapping from the first step of the analysis and the accuracy from the previous step, we can also find if there are any accurate or partially accurate location-cycles of items where items were potentially assigned to the wrong location (i.e., within the threshold we have defined for our binary location accuracy judgement but with a center on a different item's studied location). The cycles are computed via a graph-theoretical connected-components analysis where closed loops in the graph of associations are treated as individual components (Hopcroft & Tarjan, 1973). If an item exists in a single-component system, it is assigned correctly with its target (studied) location. Otherwise, the items in the component system form a cycle, where the simplest cycle would be a 2-cycle, which can also be thought of as a swap, and higher order cycles can also be found (up to a 6-cycle, as 6 items are being tested). In the general case, an N-cycle is possible where N is the number of items under test. Note, this "swap" is computed differently than our previous version of "swaps" (Watson et al., 2013), as the metric described here includes location accuracy (with item identity ignored) as part of the requirement for being called a "swap". We believe that this new metric is more consistent with the general item-location misbinding (in this case also called an identity-location misassignment) interpretation of a "swap" error. These error types are collectively called compound errors.

In summary, this process produces a set of nine output metrics. The 1&2) translation (x and y), 3) scaling, and 4) rotation magnitudes from the transformation step, 5) the number of items within the location accuracy threshold after removing both the item identities and the number of compound errors of various types. For consistency with our previous work, we will consider 6) 2-cycles, aka swaps, in isolation of 7) >2-cycles in statistical analysis. Further, any system with one or more inaccurate items is treated as a 8&9) "partial" cycle (or swap), versus those in which all locations contain an accurately placed item.
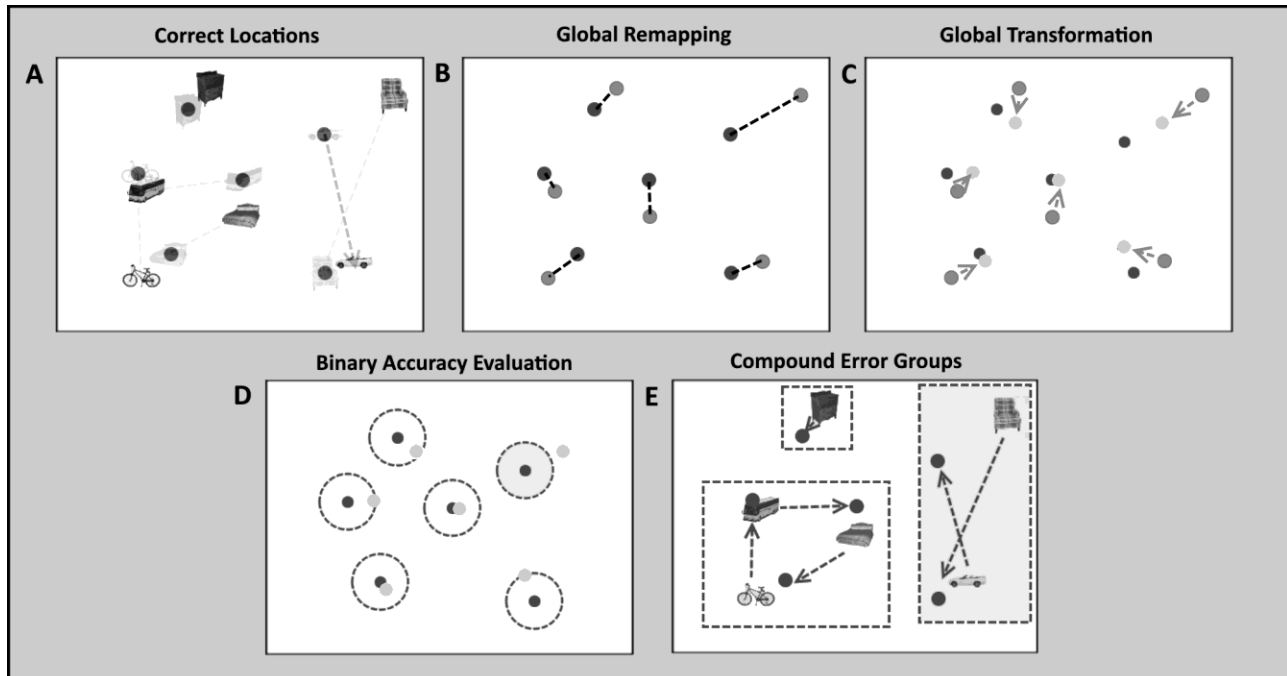
Figure 2.3: The analysis procedure as defined in the above sections. (A) shows the correct locations (dots with shadowed identities) and the reconstructed locations (items) along with shadowded dotted lines showing the original misplacement. Note that large amounts of misplacement error is accumulated via improper assignment of identities to locations. (B) shows the global remapping (dark dots are targets and light dots are reconstructed points) wherein identity is stripped away and a minimal one-to-one mapping is formed. (C) shows the global transformation (in this case a slight rotation and scaling error) where dark dots are targets, medium dots are reconstructed points, and light dots are post-transformation reconstructed points (i.e. points where the global error have been subtracted from the reconstructed points). (D) shows the binary accuracy evaluation on the post-transformed points (where the dark shaded circle signifies an inaccurate item). (E) shows the compound errors identified and labelled according to the previous steps (where shaded regions are considered "partial" groups where at least one item was inaccurately placed).

2.2.3.5 Gestalten Shape-Like Information

There are many ways to define "shape", and this proposal evaluates only one such definition. Shape is defined here via the location vertices and their relative positions, such that the shape is altered with item placements in non-studied locations and unaltered with transformations such as translation and scaling. Under this definition of shape, shape can be thought of as the highest order of compound identity-location and item-item relational information which includes all pairwise relations amongst items without regard for item-environment relations. One way to dissociate identity-location and item-item relations is to strip away the identity information and treat all items as if they were anonymous (i.e. unlabeled) in order to determine if the placement locations match up with any studied locations (see Identity Remapping above for details). The current study takes steps to evaluate and dissociate the

degree to which shape-like information may be maintained by hippocampal patients by comparing the number of correct location placements (disregarding item identity) and the magnitude of different types of global errors (i.e. error which is shared between multiple items in the reconstruction), such as translation (i.e. a fixed offset in x and/or y of all items), scaling (i.e. all items being closer to or farther from a central point without a change in angle between any item pair), and rotation (i.e. all items rotated some distance around a central point) between the groups which might suggest differences in shape memory.

2.2.3.6 Statistical Analysis

All t-tests shown were conducted via Welch's t-test for unequal variance and unequal sample size (two-tailed) unless otherwise indicated. An alpha level of 0.05 was used to determine significance. This means that the degrees of freedom for each t-test is the adjusted degrees of freedom given by the Welch-Satterthwaite equation which corrects for unequal sample size and variance. Repeated measures ANOVA were used when comparing metric performances and set size data. All ANOVA used Huynh-Feldt correction for non-sphericity.

2.2.3.7 Reanalysis of Watson et al. (2013) Set Size Data

Previous work from our group showed that hippocampal patients committed swap-like relational memory errors even at set sizes as small as 2 items, whereas such errors were virtually non-existent in comparison participants even at size sizes of 5. These data suggest that hippocampal damage leads to a pervasive inability to retain relational information even in the most restrictive circumstances. To investigate whether all types of relations are set size invariant, we conducted a reanalysis of a previously acquired data set (Watson et al., 2013) to determine if the patterns seen in the new data set hold in the previous one. The data contained four patients and four matched comparisons. Because that data set contained variations in set size, a repeated measures ANOVA was run to determine if the average error magnitudes across set sizes differed for each type of error. There was an overlap of one patient in both studies, 1846, though the data were collected approximately 4 years apart.

2.3 Results

2.3.1 Differences in Misplacement Accounted for by Global Errors

Consistent with previous work (Jeneson et al., 2010; Smith & Milner, 1981; Watson et al., 2013), the hippocampal patients show significantly more overall misplacement (difference between means is 109.2 pixels; $t(4.85) = 3.96, p = 0.02$) than the comparison participants (**Figure 2.4** "Original"). However, our new analysis framework allows us to determine the degree to which this misplacement difference can be accounted for by the different types of spatial relational errors. The first step in our analysis process

was to remap the item-location relations to quantify the misplacement accounted for by identity-location errors. This quantity is the sum of all of the misplacement accounted for by each identity-location error (**Figure 2.4A** "Post-Remapped") and, indeed, performing this procedure removes the majority of the difference in misplacement between the two groups (difference between means is 26.6 pixels; $t(9.33) = 2.59, p = 0.03$). This procedure also shows a significant difference in the magnitude of the reduction (**Figure 2.4B** "Remapped Error") of overall misplacement between patients and comparison participants ($t(5.67) = 2.81, p = 0.04$) as would be expected given the relative magnitudes of the original misplacement. Note that the reduction in misplacement from remapping accounts for over 75% of the difference between the groups. In other words, the majority of the difference in performance between patients and comparison is due to identity-location errors.
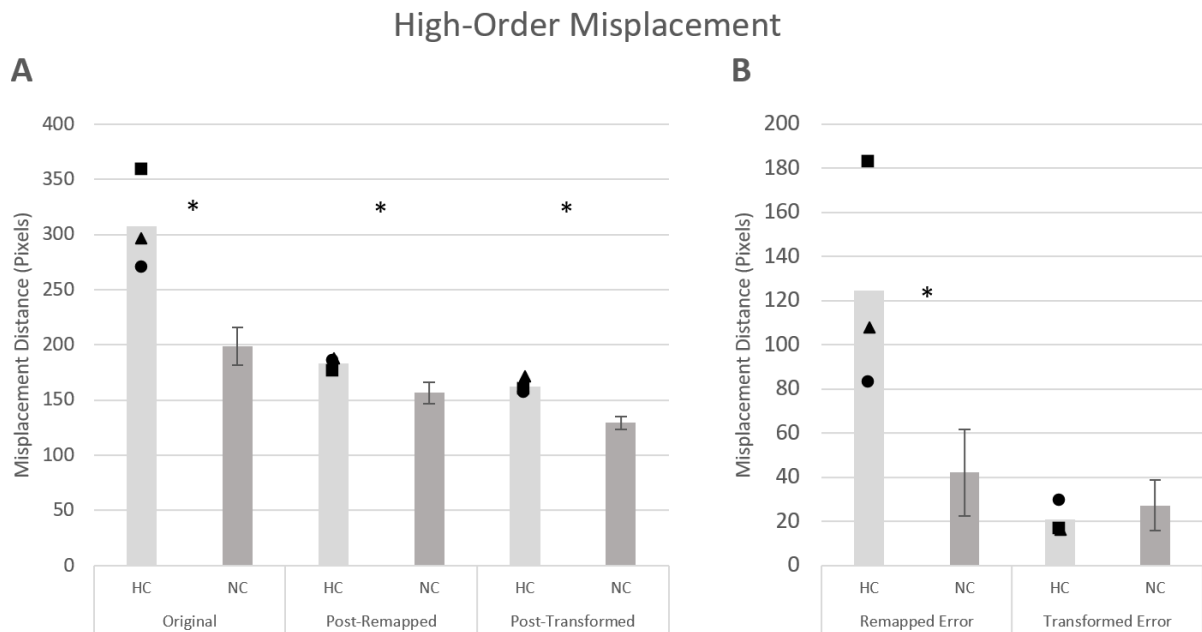

Figure 2.4: Comparisons of Patient (HC) vs Comparison Participant (NC) groups for A) Misplacement Across Analysis Steps, i.e. misplacement magnitudes at each step of analysis and B) Reduction in Misplacement by High-Order Error Type, i.e. differences in misplacement accounted for by each step (i.e. the subtraction of A). For Patients, individual symbols are shown such that: Square is 1951, Triangle is 1846, and Circle is 2563. (* means *p*<0.05)

Regarding the transformation misplacement (**Figure 2.4A** "Post-Transformed"), a repeated measures ANOVA was performed (Huynh -Feldt corrected) with group by misplacement (original, post-remapped, and post-transformed). A main effect of both group (*F*(1, 10) = 8.097, *p*=0.02) and misplacement (*F*(1.34, 13.41)=45.392, *p*<0.001) was found, as well as an interaction between group and misplacement (*F*(1.34, 13.41) = 7.536, *p*=0.01). Follow up t-tests confirmed both patients and comparison participants show a

significant reduction in their misplacement via the transformation procedure ($t(3.82) = 4.18, p = 0.03$ and $t(13.41) = 2.36, p = 0.03$ respectively). However, the patients and comparison participants show no difference in the amount of misplacement accounted for by the global error amongst their items ($t(9.89) = 0.51, p = 0.62$). When the different transformation components (translation, scaling, and rotation; **Figure 2.5**) are observed independently, we see that only rotation has a significantly different magnitude, with patients having larger magnitude rotation errors than comparisons ($t(9.18) = 2.72, p = 0.02$).

Once all global transformation procedures are performed, patients and comparisons still show differences in the remaining misplacement (difference between means is 33.0 pixels; $t(9.54) = 4.57, p = 0.001$). This misplacement is "local" misplacement, which is not shared by all items nor accounted for by identity-location errors. Note (**Figure 2.4A** "Post-Transformed") that because there is no significant difference between patients and comparisons in the amount of global misplacement error (**Figure 2.4B** "Transformed Error"), the final difference between the groups is still significant.

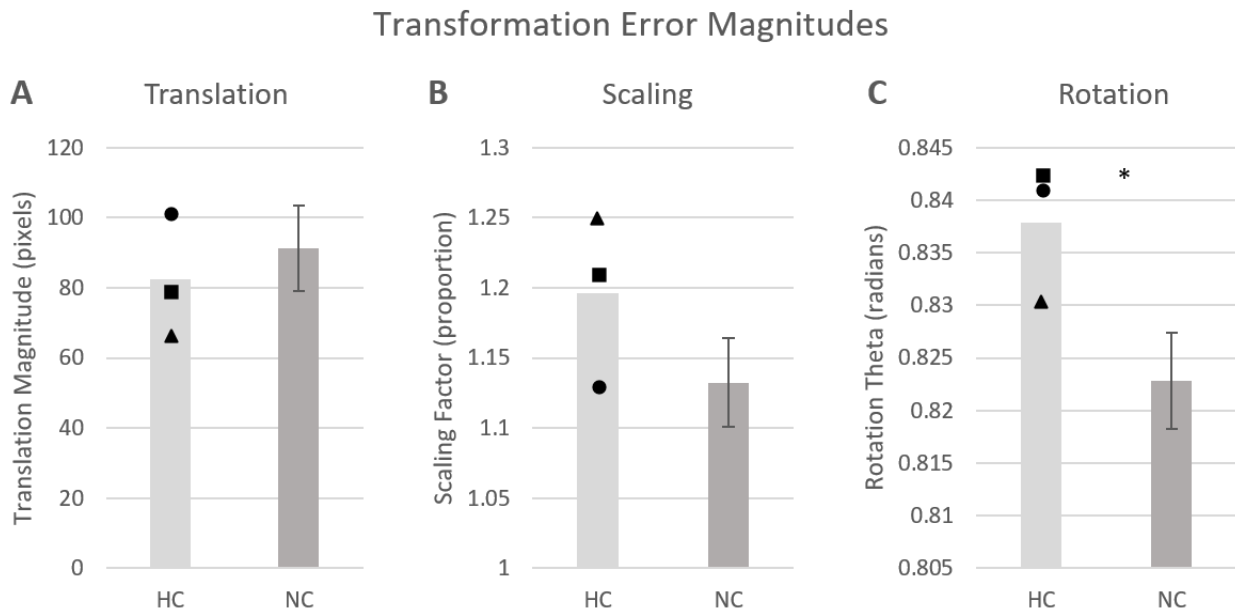**Figure 2.5**: Comparisons of Patient (HC) vs Comparison Participant (NC) in the magnitude of the three types of global transformation error; A) the Translation magnitude, B) the Scaling magnitude (where 1 means no scaling, <1 means shrinking, and >1 means stretching), and C) the Rotation angle (in radians). For Patients, individual symbols are shown such that: Square is 1951, Triangle is 1846, and Circle is 2563. (* means *p*<0.05)

2.3.2 First-Order and Higher-Order Relational Errors

When we evaluated the location placement accuracy of identities to their assigned locations, there were three primary first-order comparisons of interest. The first two comparisons addressed the question of whether patients could remember identity-location information as often as comparisons, evaluated by: 1) the number of correct placements and 2) the number of items placed in another item's location. The third comparison of interest addressed the question of whether the patients could remember just the location information as often as comparisons regardless of identity, which was evaluated by examining the number of items placed in any valid location (note that this is just the sum of 1 and 2). To examine group differences in these quantities (particularly 2 and 3), we performed a between groups repeated measure ANOVA on the identity-location errors (i.e. 2, above) and the total number of items placed in any valid location (i.e. 3, above, the sum of the identity-location errors and the correctly placed items). We found a significant interaction of group and these two quantities ($F(1, 10) = 17.127$, $p=0.002$). Follow up Welch's t-tests showed that hippocampal patients made fewer accurate placements of the correct identity to the correct studied location (e.g., successful identity-location binding) than comparison participants ($t(5.21)=5.45$, $p=0.003$; **Figure 2.6**). Importantly, however, patients had more placements of items in another item's studied location (i.e. correct item-location relations with incorrect identity-location relation) than comparison participants ($t(5.06)=2.83$, $p=0.04$; **Figure 2.6**), even after subtracting the global errors. When we looked at memory for locations (regardless of identity; i.e. items placed in any valid location), we found no significant difference between patients and comparisons ($t(4.71)=0.71$, $p=0.51$)).

In regard to compound errors (i.e. swaps or cycles), we saw no difference between patients and comparison participants ($t(9.96)=0.39$, $p=0.71$ for swaps and $t(7.11)=1.75$, $p=0.12$ for cycles, respectively) despite large differences in individual identity-location assignment.

The fact that patients and healthy comparison participants had similar memory for locations (ignoring identity assignment) serves as the first piece of evidence suggesting the possibility that some shape-like information (albeit only in the form of a set of vertices) may be used by patients to reconstruct spatial location. In any case, it is clear that some aspects of relational information in the reconstruction (i.e. some combination of item-environment or item-item relations) was not impaired in hippocampal damaged patients.
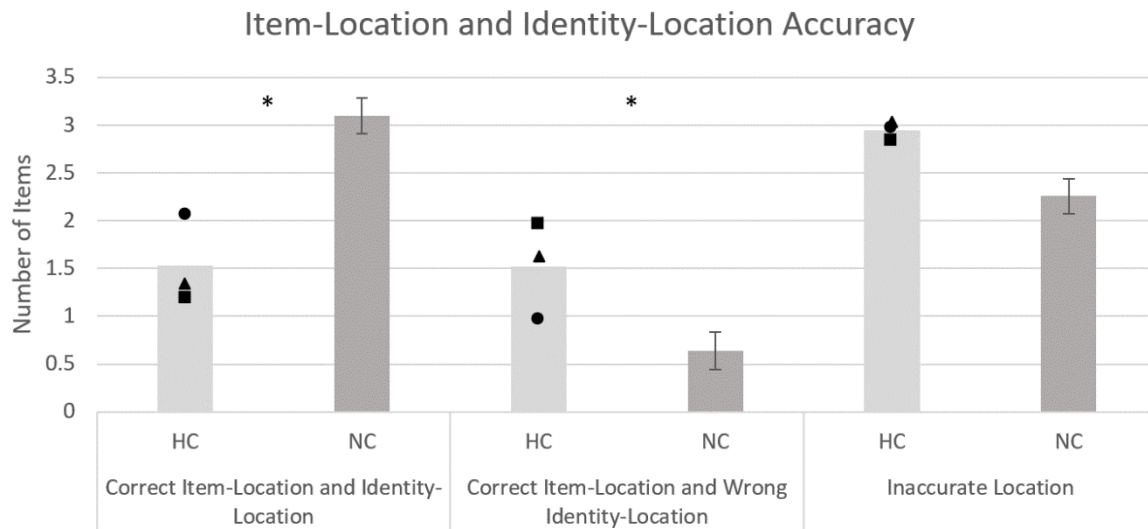
Figure 2.6: Comparisons of Patient (HC) vs Comparison Participant (NC) in the number of items in which item-location and identity-location was correct and incorrect. Note that items which were not placed in any studied location are collapsed together as 'Inaccurate Location'. For Patients, individual symbols are shown such that: Square is 1951, Triangle is 1846, and Circle is 2563. (* means *p<0.05*)

2.3.3 Differences in Accuracy of Item-Location Associations across Set Sizes

As our previous work showed that patients committed relational memory errors even at very low item set sizes, it is useful to ask if, given our new framework for analyzing various types of relational information, we see the same pattern of impairment across all set sizes. We took advantage of our previous data set from Watson et al 2013 to determine if the number of studied items impacts the different types of spatial-relational errors discussed here. In a 2 x 4 ANOVA with factors group (patients and comparisons) and set size (2, 3, 4, 5), hippocampal patients showed no main effect of set size $(F(3,9) = 0.746, p = 0.552)$ on accuracy of placing items in their target locations, placing roughly only a single item in its correct location (identity-location relation) regardless of set size (2 to 5). However, the main effect of set size was significant for comparisons $(F(2.64, 7.93) = 4.51, p = 0.04)$. Importantly, as reported previously (Watson et al., 2013), only patients committed identity-location errors at small set sizes; comparisons did not. Furthermore, when considering performance at set size 5, in which both groups made identity-location errors, there were no significant group differences for other types of compound relational errors, namely swaps and cycles; swaps $(t(5.0) = 0.2597, p = 0.8055)$ and cycles $(t(5) = 0.1800, p = 0.8642)$. Taken together, these data show that across set sizes, patients show specific impairment to identity-location information without showing any significant

compound relational error differences once the set size is sufficiently high for comparisons to begin
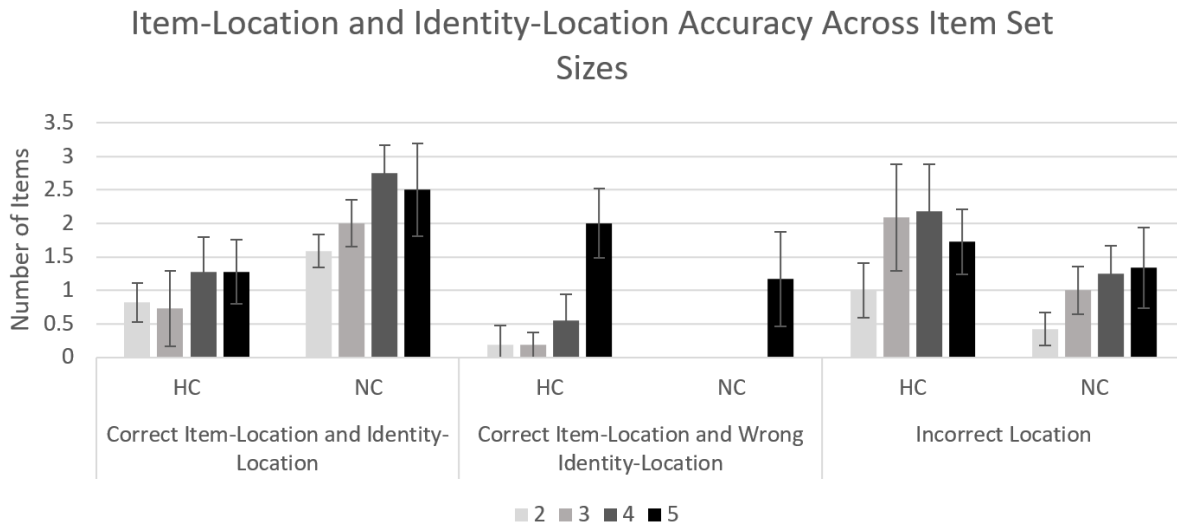
making identity-location errors.

Item-Location and Identity-Location Accuracy Across Item Set Sizes

Figure 2.7: Comparisons of Patient (HC) vs Comparison Participant (NC) Item-location and identity-location accuracy across set sizes for patients vs comparisons. Note that items which were not placed in any studied location are collapsed together as 'Inaccurate Location'.

2.4 Discussion

The present study investigated the impact of hippocampal damage on various types of relations in a spatial memory task. While the role of hippocampus in spatial memory has an extensive history (Hayes et al., 2004; O'Keefe & Nadel, 1978; Ryan et al., 2009), the present data demonstrated that patients with hippocampal damage are not impaired relative to healthy comparison participants on *all types* of spatial errors; rather they demonstrate particular deficits in identity-location relational memory.

There are many reasons why identity-location relationships, specifically, might be impaired with hippocampal damage. Consider a case where a single item is presented in a bounded spatial environment. Only one type of relational information, namely the item-environment relation, is necessary to remember where that item belongs. Remembering the identity of that item is not necessary to complete the task. However, if a second item is added to the display, each item has an item-environment relationship, but also item-item relations and identity-location relations. It is, of course, not necessary to encode all three relation types to represent the display because some relations provide redundant information (i.e. the item-item relations contain information which is redundant with the item-environment relations). However, for an accurate representation, the identity-location relationship is not redundant with other types of relations and must always be considered. Identity-

location relations are entirely arbitrary (in the SR task) and are the critical component for successful memory performance; perhaps unsurprisingly then, arbitrary identity-location relationships are the type of relations that are most hippocampally dependent in a spatial memory task. As mentioned previously, it is important to note that item-item and item-environment information in this task and framework are not necessarily equivalent to seemingly similar relations in other domains (i.e. item-item where both are words or item-environment where it is the mapping of an object *including its identity* to a spatial location); this is due primarily to the specificity with which we have defined these terms (i.e. 'items,' in this case, have had their identity information removed and specifically reference spatial-only information).

The assertion that the *arbitrary* type of relational information in this task is the type showing impairment is consistent with previous findings showing hippocampal damage impairs memory for *arbitrary* relations between words (Giovanello et al., 2003), spatial locations in a scene (Hannula et al., 2006), and temporal relations, and, indeed, all manner of relations regardless of the informational domain (Konkel et al., 2008). The data bear out this conclusion by showing that if we remove the identity information from the patient data, we see the difference between the groups in overall misplacement reduced by over 75%. Furthermore, when stripped of identity information, patients and healthy comparison participants do not differ on the number of correctly chosen locations. This shows that although patients show impaired memory for arbitrary relations in all stimulus domains, including spatial relations, they are not impaired on all *types* of spatial relations (i.e. they are able to reconstruct some item-environment and item-item relations). Patients with hippocampal damage show a deficit in reconstructing the *specific* item identity to its *specific* studied location.

The present data also demonstrated that hippocampal patients and comparison participants do not differ in compound spatial relational memory errors (e.g., translation, scaling, swapping, and cycles). And while rotation errors showed a statistical difference in magnitude between patients and comparison participants, the magnitude was quite small (0.86 degrees), accounting for very little overall misplacement difference between the groups. Moreover, because the very first step in the analysis is to remove identity information, the total amount of possible rotation is artificially restricted to no more than the half of the maximum angle between two items from the rotation center, making it more complicated to interpret global error type than scaling and translation that have no such restrictions. It is, of course, possible that the significant difference in rotation errors is due to additional memory deficits. Unlike scaling and translation, in which pairwise spatial relations are invariant (i.e. if item A is

above and to the right of item B and a translation or scaling transform is performed, item A will always still be above and to the right of item B), pairwise spatial relations are not invariant under rotation. This could point to further, more subtle deficits in particular types of relations, but further investigation into this error type is needed.

It is interesting to note that some misplacement remains as local, "noisy" misplacement after the removal of identity and global errors, and this local misplacement was significantly greater in patients than comparison participants. This is consistent with other accounts of spatial memory and the hippocampus such as the Precision and Binding Model (Yonelinas, 2013) which posits that the hippocampus is critical for "high-resolution" binding in space, predicting that hippocampal patients should have poorer spatial memory precision. These accounts may not generalize to all domains of information, however, as some evidence shows hippocampal damage relating to impairment in the quantity of information which can be remembered, rather than the quality (or precision) of the information when testing hippocampal patients on color information (Warren et al., 2015).

In regard to compound relations involving identity-location information (e.g., swaps and cycles), it has been proposed previously that a lack of memory for identity-identity and identity-location relations in hippocampal patients may result in more identity-identity swapping (Watson et al., 2013). However, the present data demonstrated no group difference in swaps or cycles after accounting for overall location accuracy. The reason for this discrepancy is likely a combination of: 1) a difference in mathematical formulation of swapping (i.e., not using accuracy as a prerequisite for a successful swap results in an over estimate of the number of swaps) and 2) the inability of hippocampal patients to represent a sufficient amount of relational information to commit compound-relational error such swaps (which requires memory for two locations and their positions relative to one another and/or the environment but no information about item-location relations) more than by chance. Specifically, on the second point, if hippocampal patients are not maintaining identity-location information, their ability to have group-wise reassignment errors will be limited because they do not have specific knowledge of the constituents of the group. These compound swap/cycle errors do occur in both patients and controls (although patients uniquely commit them at very small set sizes), but once the set size is sufficiently high (i.e. once the comparison group actually makes identity-location errors) the groups do not differ in the frequency of their occurrence. Said differently, just as identity-location errors seem to account for a substantial proportion of patients' overall misplacement errors, the disproportionate occurrence of

39

swap and cycle errors at small set sizes in patients can also be attributed to their overall deficit in identity-location binding.

What causes this sudden increase in identity-location errors in comparison participants at higher set sizes? One possibility requires consideration of accurate identity-location placements in isolation of other metrics. Across all set sizes, hippocampal patients make approximately one accurate identity-location placement. This accuracy could be due, in part, to a maintenance of a single identity-location relation in working memory across the relatively short time span between study and test. Comparison participants, on the other hand, make accurate identity-location placements proportional to the set size.

When we consider how hippocampal patients may be able to maintain item-environment and item-item spatial relational information, it is interesting to consider alternate representations which may be able to maintain those relations. Previous work suggests that hippocampal damage may not impair shape information in SR task performance, as hippocampal patients may be able to represent shapes via an alternative, "unitized" representation, which can be formed via Gestalten perceptual features, i.e. configural features which constitute parts of a unified whole, rather than relations (Corsi, 1972; Kessels et al., 2000; Uttal & Chiong, 2004). The present data provide a suggestive look at this assertion from a purely vertex-oriented definition of shape, as the patients and comparisons did not differ in the number of locations (i.e. vertices) in which they accurately placed an item (independent of item identity), nor did they differ in the amount of global misplacement (i.e. misplacement that is systematically shared by all items) of the entire item set. This suggests that, at least in this highly restrictive definition of shape being defined via vertices, a similar amount of shape-like information is present in both comparison participants' and hippocampal patients' reconstructions. That said, using global misplacement transformations as a measure of shape is not necessarily the optimal method for investigating this issue, even with shape being defined solely via vertices. Topological Data Analysis techniques may prove to be a better analysis framework for determining the degree to which high-level relational information is preserved in a reconstruction (Carlsson, 2009; Ghrist, 2007). In Topological Data Analysis (TDA), large groups of vertex relations can be evaluated and compared directly with minimal assumptions by forming a more mathematically rigorous definition of the high-order structure of these relations. A similar technique to TDA has been used previously in spatial memory by Kolarik et al. in analysis of data from hippocampal damaged patients using a virtual Morris Water Maze (Kolarik et al., 2016, 2017). It is unclear, however, if this technique scales well to a very small number of points, and, as such, these

analyses are not included in this work; future investigations better suited to the technique should be considered.

There are alternate explanations for the ability to maintain item-environment and item-item spatial relations other than maintenance of shape-like information. It is possible individual relational information *is* being stored separately via a different brain region(s). For example, the parahippocampal place area has been shown to be involved in processing geometric information about scenes (Epstein, 2014), and the occipital place area has been related to processing of boundary relations between items and visible scene boundaries (Julian et al., 2016). If shape information were being used, regions such as the inferior temporal cortex could be involved in providing preprocessed representations of shape via shape selective neurons identified in this region (Perrett & Oram, 1993). Alternatively, some heuristic approach to reconstruction could be used which takes advantage of hidden constraints in the task (i.e. it is unlikely for two items to appear on top of each other, which constrains the possible reconstructions). Should the hypothesis that shape-like information is being used to maintain these relations be investigated, different experimental paradigms could be used in which various levels of complexity of shape-like information are superimposed on a point set and presented to participants to be learned then reconstructed. In particular, edges between the locations could be drawn to avoid the need to infer some item-item relations which may contribute to a shape-like representation. We would predict that patients would perform similarly to comparisons for simple shapes and show deficits primarily in identity information for the vertices. Patients may also show performance deficits for more complex shapes, which are not as easily unitized (thus preventing a non-hippocampal representation from being as helpful).

Taken together, the evidence presented here defines the various *types* of relations present in a spatial-relational paradigm by applying a systematic framework that allows a more thorough investigation of the role of hippocampus in spatial-relational memory. Consistent with previous theories of the role of hippocampus in relational memory via the arbitrary binding of information (Cohen & Eichenbaum, 1993; Eichenbaum & Cohen, 2001), patients with hippocampal damage demonstrated specific deficits in arbitrary binding of identity-location relational information. Patients do not show a deficit in compound relational errors or location accuracy, suggesting that some types of relational information (in fact, the non-*arbitrary* relational information; i.e. information which has alternative derivations or redundancy) can be maintained by memory systems which are not hippocampally dependent. These results further clarify the role of the hippocampus in relational memory theory to be specifically in maintaining

arbitrary relations which cannot be derived from other information which might be remembered by the participant. We have also proposed several possible alternate means through which other, non-arbitrary information may be maintained (such as a shape-like representation), which can be investigated in future research. We would predict, however, that because one benefit of representing arbitrary relations is in the ability to flexibly use those relations with changing task demands, whatever alternative means is being used to maintain these relations may not permit the same amount of flexibility.

SR tasks have proven enormously informative in that they provide sensitive measures of relational memory abilities that have been tied not only to hippocampal damage, but to more subtle hippocampal deficits such as subjective memory complaints (Lucas et al., 2016), spatial pattern separation (Clark et al., 2017), and medial temporal lobe viscoelasticity (Schwarb et al., 2017, 2016). Precision and specificity of the metrics and the ability to distinguish between several different types of spatial relational errors will undoubtedly serve these lines of investigation well as it will allow a far more resolute picture of the specific impairments that may be tied to particular behavioral, functional, and structural phenotypes.

Chapter 3: Memory during Time Travel: Spatiotemporal Navigation, Contextual Boundaries, and Relational Memory Errors in Virtual Reality

3.1 Introduction

Episodic memory involves memory for events, i.e. things which happen in time and space, organized according to relational and contextual information, and is often referred to as the "ability to jump back in time and relive events from the past" (Tulving, 2002). In humans, the *hippocampus* has been shown to be a critical brain region for episodic memory such that hippocampal damage in humans causes pervasive *amnesia* (for review, see Cohen & Eichenbaum, 1993). There has been considerable historical success in studying hippocampal function in the spatial domain, with early experiments identifying "place cells" in the hippocampus of rodents which fire in association with the location of an animal in a spatial environment (O'Keefe & Nadel, 1978). The temporal organization of memory has also received considerable attention, including the recent identification of "time cells" in the hippocampus which fire selectively for moments in time (Eichenbaum, 2014). However, experiences are not restricted to space or time in exclusion of each other, but consistently involve a mix of both. Despite this, most studies do not simultaneously examine behavior in both space and time, possibly because space, as an explorable domain that can be selectively sampled, and time, as unidirectional flow of experience that cannot be selectively sampled, seem to be too distinct for direct comparison. This work takes the first step towards resolving this asymmetry by allowing both space and time to be explored by participants using *Virtual Reality* (VR), a technology which allows for increased measurement and control fidelity as well as increased ecological validity when compared to other methods (Schultheis et al., 2002). This approach provides a new framework to investigate and understand the nature of episodic memory, extending beyond the limitations of typical experimental environments.

Relational memory theory explores the relationship between different domains of information (such as, but not limited to, space and time) and their representation in the hippocampus. (Cohen & Eichenbaum, 1993; Eichenbaum & Cohen, 2001; Konkel & Cohen, 2009; Konkel et al., 2008). According to relational memory theory, the hippocampal region critically supports episodic memory via relational memory, i.e. the learning of arbitrary relational information regardless of the domain (Konkel et al., 2008) – a function which is pervasively impaired in hippocampal damaged patients. To evaluate relational memory in humans, *spatial reconstruction* (SR) tasks have proven particularly effective in examining relational memory errors in space. This task can be used to infer an individual's success in forming arbitrary associations between items and locations in space (Huttenlocher & Presson, 1979), and to determine the susceptibility of making relational memory errors due to hippocampal damage (Jeneson et al., 2010;

Smith & Milner, 1981; Watson et al., 2013). In an SR task, participants are asked to study the locations of items in space and, after a short delay during which the items are removed from the environment, place the items back where they were studied (i.e. reconstruct the study configuration). Hippocampal damage specifically impairs arbitrary identity-location binding, i.e. binding of the identity information for a specific item (i.e. the item was a boat, plane, boxing gloves, or flower) to a specific location (Horecka et al., 2017, under review). In healthy individuals, identity-location binding errors and compound memory errors (i.e. those in which groups of items are misbound to each other's locations) can occur when the number of to-be-remembered items increases beyond approximately 4-5 items (Watson et al., 2013). SR tasks have been used successfully in examining the relationship between relational memory and subjective memory impairment (Lucas et al., 2016), medial temporal lobe viscoelasticity (Schwarb et al., 2017, 2016), and spatial pattern separation (Clark et al., 2017), however, time has always been intentionally isolated into other tasks. In this work, we present the first S*patiotemporal Navigation Task* in which the locations of events in both space and time are studied and reconstructed.

There is substantial evidence that the hippocampus is critical for organizing information about the temporal order of events, and that in humans, hippocampal damage results in deficits in temporal-order memory for those events (Dede et al., 2016). Additionally, functional imaging studies have shown hippocampal activation during encoding and retrieval of event order information (see Eichenbaum, 2014 for a review). In rodents, hippocampal lesions also impair temporal-order memory for non-spatial stimuli without impairing memory for the specific stimuli (Fortin et al., 2002; Kesner et al., 2002). However, in all of these examples, event series have been viewed unidirectionally, critically restricting the directionality of temporal relations and information. Although physical space-time may have the restriction of direction-of-flow in time, this restriction can be, at least in part, artificially overcome in simulation.

Additionally, there is extensive evidence that contextual information influences temporal order memory in online processing of sequences (Ezzyat & Davachi, 2014). For example, Zacks et al. (2007) have suggested that an *event model* which is updated according to event *boundaries* can describe processing of events in time. In humans, *context boundary effects* (i.e., distance judgements of item pairs within vs across context boundaries are biased such that items which are separated by a context boundary are judged as further apart than those which are within a context) have been demonstrated (DuBrow & Davachi, 2013) and the hippocampus has been implicated in representing temporal order and distance in the presence of changes in context (Davachi & DuBrow, 2015). Taken together, these studies highlight

the importance of context in episodic memory representations; importantly, however, this evidence was bound to a unidirectional timeline of events and often discretized such that time is not allowed to flow normally in a continuum (with the exception of the Zacks work), but is instead broken up into discrete events. It is unclear if, given the ability to freely explore a series of events organized by context in a continuous timeline, these same effects would persist. In particular, the ability to cross contextual boundaries in both directions as often as desired (a feature of real-world navigation which is far more common than the restriction that each boundary may only be experience once and in the same direction) has never been explored.

In this work, we introduce a powerful *Spatiotemporal Navigation Task* for examining relational memory errors and the impact of context on memory for event locations in space *and* time by allowing participants to freely explore an environment in both space and time using Virtual Reality. Critically, we provide participants with the ability to "time travel" (i.e. decide the direction of the flow of time), placing the ability to explore space and time on more equal footing so that errors in either domain can be examined simultaneously. This "time travel" ability is necessary to allow spatiotemporal navigation where participants can visit events in time as freely as they can the spatial locations of events. To our knowledge, this is the first *Spatiotemporal Navigation Task*.

3.2 Methods
3.2.1 Participants
43 participants (ages 18-24; mean 20.0 years; 23 males; 13.8 mean years of education) from the University of Illinois Urbana-Champaign participated in the study. In addition to the 43 participants being analyzed, 4 participants were excluded due to not completing the task within the allotted 2.5 hours. All participants signed informed consent prior to participation, approved by the University of Illinois Urbana-Champaign Institutional Review Board and were paid $8 per hour of their time.

3.2.2 Design and Procedures
In this task, participants were placed in a virtual environment measuring 40 meters by 40 meters, bounded by walls on the exterior (see **Figure 3.1A)**. The environment resembles a small park with trees, rocks and plants. Participants used the Oculus Rift DK2 and a wireless Xbox controller to interact with the environment. Auditory cues were presented via a pair of over-ear, stereo headphones. Participants were given a single practice trial to learn the simulation mechanics as well as to gain competency in using the virtual reality hardware. Practice was not timed and participants were allowed as much time as they desired to become comfortable. The stimuli during practice were white boxes (0.5m x 0.5m x

0.5m) floating 0.25m above the ground with gray dots on each face. Five items were used in practice. In order to encourage exploration of the environment, an invisibility bubble 10m in diameter was placed around the participant such that any items beyond that distance away from the participant were not visible. A ring was visible on the ground to indicate to the participant the distance at which items were visible. In addition, some items were always present throughout the trial, while others underwent events (i.e., entered the display or were removed from the display) at specific time points during the simulation timeline (see **Figure 3.1** for a description). Thus, there were three event types 1) items flew into the sky and left the environment at specific times (an "up" event, i.e. disappearing), 2) items fell from the sky and land in the environment at specific times (a "down" event, i.e. appearing), and 3) items were stationary and thus always present in the environment. During practice, there were two up events, two down events, and one stationary item. An auditory cue was given every time an event occurred regardless of whether or not the item was in view. Item positions were determined via random number generator; however, they were constrained in time to balance for within vs. across context boundary distance.

Temporal context information was provided in the form of the simulation boundary walls changing colors at evenly spaced intervals throughout the simulation. Thus, context boundaries were defined as shifts from one color to the next. The colors yellow, red, green, and blue were used for the contextual information, which changed every 7.5 seconds during practice such that all contexts appeared within 30 seconds. The order of the temporal contexts was counterbalanced across participants with half receiving yellow, red, green, and blue and the other half receiving blue, green, red, and yellow.

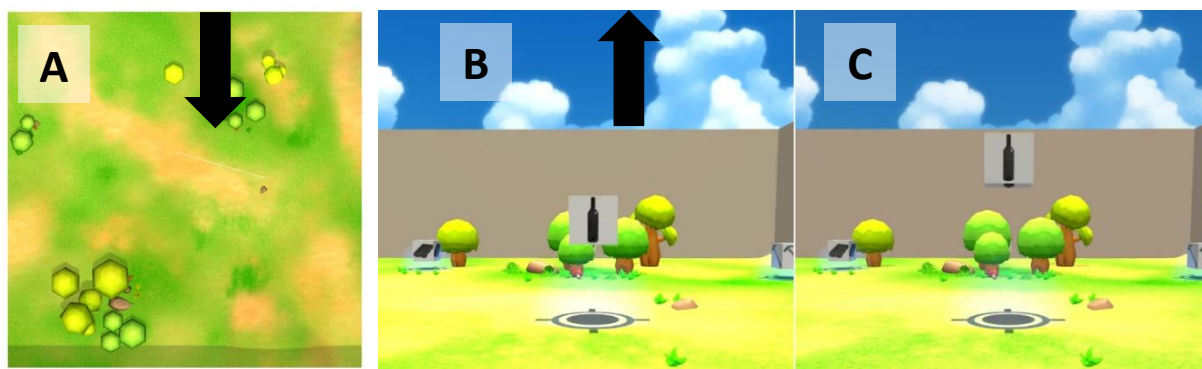**Figure 3.1**: The two types of events an item can undergo in time and a visualization of the environment. A) shows a "down" event, in which the item falls from the sky, B) shows an "up" event, in which the item flies into the sky, and C) shows the environment from a top-down view. Note that these two types of events are time-symmetric, in that they look identical when viewed from opposing directions (forward or backward) in time.

Participants could both walk (via pressing "forward" on a joystick) around the environment (and turn their head to look in any desired direction). Temporal navigation was controlled via button press that could reverse the flow of time, at the same speed as the forward flow of time. To provide additional temporal information, ambient music was played for the entire timeline of the task, meaning reversing the flow of time reversed the music as well. This music, however, provided very little information beyond an auditory sensation of the passage of time as it did not contain any significant changes in style, tone, or mood through the one-minute timeline. Participants' looking direction and location in space and time were logged at every time point in the simulation (~16ms intervals). If the participant reached the beginning or end of the simulation's time course, time was paused but the participant was still allowed to freely move about the spatial axes in the environment. Participants were instructed to find the locations of each item in space and time, specifically when their events happened. To signify they had found the item at its event, there was a 0.5 second window during which participants could see the item's event (if close enough) and press a button on the controller to acknowledge they witnessed the event. When the participants successfully click on the item during its event, the item would change to the background color during its event (with stationary items staying grayscale when clicked). Participants were given unlimited time to explore the environment, find each item, and acknowledge they had viewed each event.

After the practice trial, there were four experimental trials each with a study and test phase. Experimental study trials were identical to the practice trial with three exceptions. First, the number of items was increased to ten. Therefor there were four up events, four down events, and two stationary items. Second, the gray dots on the sides of each white box were replaced with namable object images (see **Figure 3.1B&C** for example). Third, the context changed every 15 seconds such that all contexts appeared within 60 seconds. Each context had two item events during its duration. Finally, the 4 study-test trials all used the same spatiotemporal item configuration, allowing learning across trials to be assessed.

During test, participants were asked to reconstruct each item in its spatiotemporal location as well as to determine which event type (up, down, stationary) the item underwent. To do this, participants were given an inventory of the items (an overlay on the screen) allowing them to freely select any item and its event type. As with study, participants could navigate space in on the x/y plane, and could navigate time via the reverse-time button. Context continued to change as it did during the study phase. Participants were given unlimited time to perform the reconstruction.

3.2.3 Analysis Metrics

Four primary analysis methods were used to evaluate reconstruction performance. First, a naïve misplacement distance error was computed across trials to confirm that participants were improving in all primary sources of information (spatial, temporal, and event type). Misplacement along continuous variables (space and time) was computed via the sum of the Euclidean distance between all placed items and their expected locations. The number of correct event types was determined by a simple comparison of the chosen event type to the expected (studied) event type for each item (i.e. summing the number they got wrong). Additionally, the participant navigation distance in space and time (separately) was observed across trials, during study to confirm that the study-time behavior was also improving. This spatial and temporal navigation distance was computed via the Euclidean distance between all points traversed by the participant in space (i.e. spatial distance in meters) and in time (i.e. temporal distance in seconds).

Second, an evaluation of relational memory efficacy was performed for identity-location misassignment in space and in time (Horecka et al., 2017, under review). Identity-location accuracy was evaluated by first stripping away the identity of the items and determining if each was placed in any item's location based on a 95% within-participant confidence interval of their misplacement. Then, identities were restored to determine if items were placed in their correctly associated locations (or in another item's location). Identity-location misassignments are the number of instances in which an item was misassigned to another item's location within a distance which was deemed "accurate" (see Horecka et al., 2017 for details). Identity-location misassignment has previously been identified as critically dependent on the hippocampus (Horecka et al. 2017). In order to isolate spatial and temporal information, two separate analyses were performed using these measures. The 1D temporal axis was analyzed (where location is defined as the position along the 1D axis), and the 2D spatial axes were analyzed (where location is defined in the traditional, spatial way).

Third, the effects of context boundaries on temporal memory distances were evaluated by normalizing the pairwise distance between within and across context items (dividing their reconstructed distance by the expected distance) then mean-centering to make 0 represent a perfect placement (i.e. d>0 represent items placed further than expected, and d<0 represent items placed closer than expected).

Finally, the impact of context on relational memory errors was evaluated by determining the number of identity-location misassignments to another within-context item (i.e. the other item with which it shared a temporal context) or an across-context item. Chance level for within- and across-context

misassignments were computed by taking the number of identity-location permutations for within and across context identity-location pairs. The two items which had no temporal context were excluded from this part of the analysis, and, therefore, the chance levels were 8/8P2 and 1-8/8P2, or 0.143 and 0.857, respectively, where nPr represents the number of r-length permutations of n items and the number 8 is due to the 4 contexts containing 2 possible within-context misassignments each (i.e. items a and b within a context can be misassigned a→b's location or b→a's location).

3.2.4 Statistical Analysis

All analyses which compare trial-by-trial performance use repeated measures analysis of variance using SPSS. Where appropriate, Huynh-Feldt correction was applied. Comparisons which involved collapsing performance across trials or comparing trial-over-trial performance used either paired-sample, two-tailed t-tests or one-sample, two-tailed t-tests (for comparison to expected chance values). All tests used an alpha value of 0.05 to determine significance.

3.3 Results

3.3.1 Overall Accuracy and Navigation (Misplacement Measures and Navigation Distance)

A 4 by 3 repeated measures ANOVAs (trial number by test domain; i.e. space, time, and event type) across trials (1-4) revealed that accuracy in both space and time as well as event type improved significantly across trials (main effect of trial F(2, 167)=252.86, p<0.001 and significant interaction between domain and trial F(6, 334)=14.56, p<0.001; see **Figure 3.2**). Follow-up t-tests between trials for space, time, and event type showed a significant decrease in misplacement in space between trials 1 and 2 (t(42)=2.90, p=0.006) and trials 2 and 3 (t(42)=3.50, p=0.001) but *not* trials 3 and 4 (t(42)=1.51, p=0.138); misplacement in time decreased significant between all trials (t(42)=7.77, 3.97, and 3.99, p<0.001 for each trial pair); number of incorrect event types decreased significantly between trials 1 and 2 (t(42)=4.61, p<0.001), between trials 2 and 3 (t(42)=5.36, p<0.001), and between trials 3 and 4 (t(42)=2.35, p=0.024); Finally, the amount of distance travelled in space and time also improved across trials (main effect of measure, i.e. spatial or temporal distance, F(1, 168)=268.27, p<0.001 and interaction of trial and measure, F(3, 168)=19.10, p<0.001). Follow-up t-tests between trials for spatial and temporal distance travelled show that spatial distance travelled decreased between all trials (t(42)=5.22, 3.07, 2.34, p<0.025 for each trial pair); temporal distance travelled (i.e. number of seconds not paused at the beginning or end of the simulation) decreased across trials (t(42)=5.86, 4.12, 3.48, p<=0.001 for each trial pair). Together, these analyses show improved navigation efficacy in addition to test-time memory performance across trials in the task.
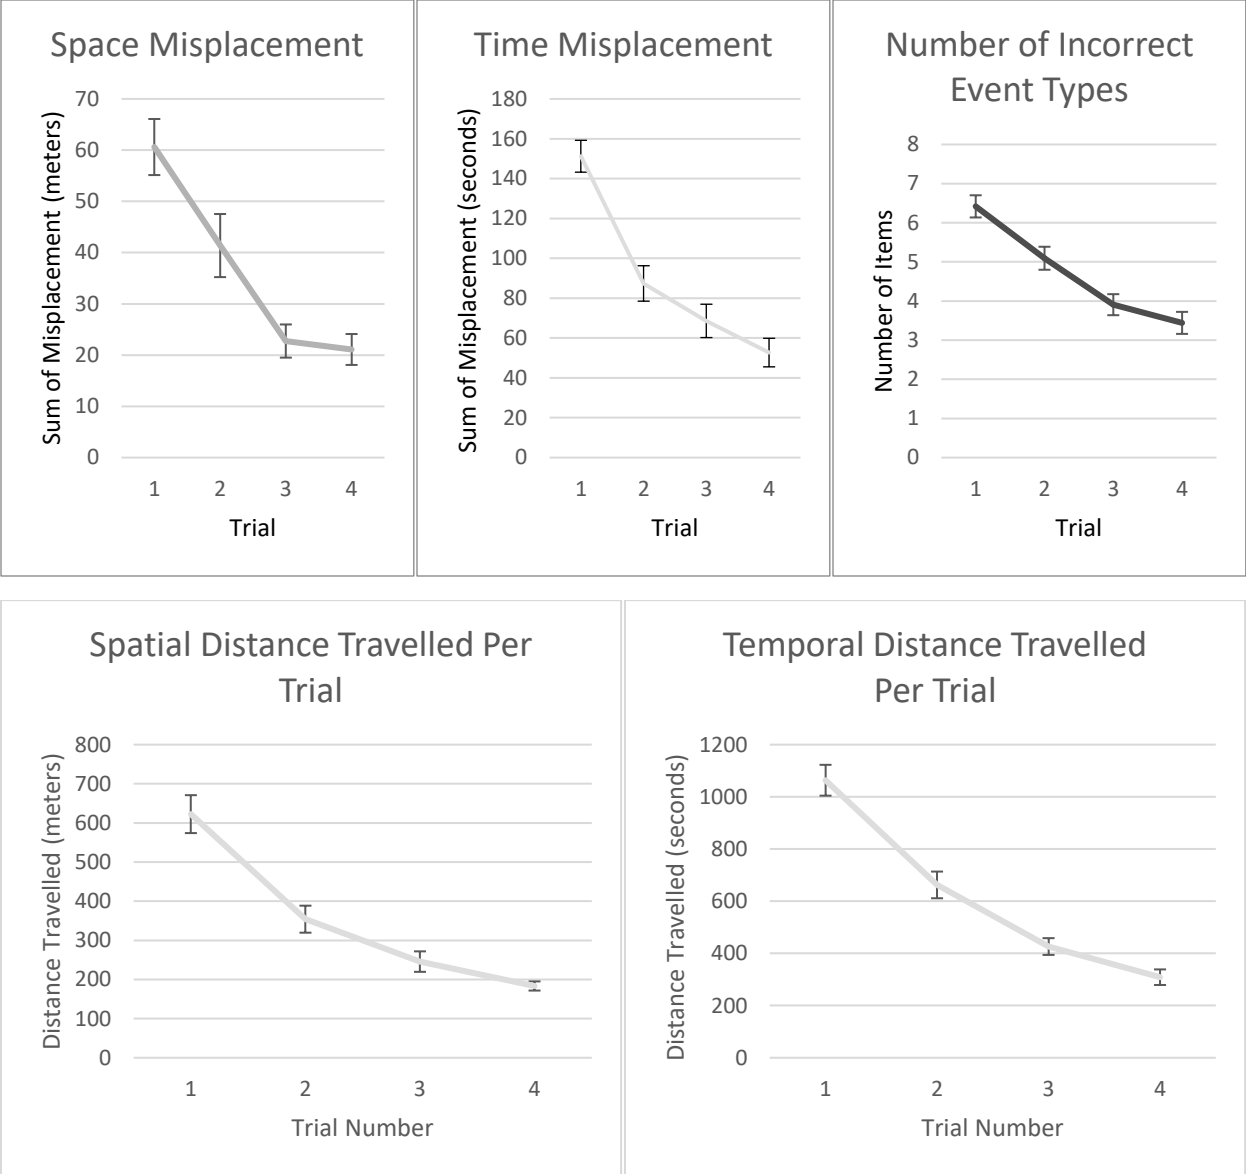
Figure 3.2: Overall errors along each of the three major information types in the task and navigation distances in space and time. Space misplacement and time misplacement are continuous variables computed via the sum of the Euclidean distance of all items in a trial to their expected locations in space and time. The number of incorrect event types is just a count of the number of items which were assigned the studied event type. Note that participants improved in all metrics between trials 1 and 4, showing the task was generally well understood. Moreover, navigation performance improved across trials in both space and time showing the improvement was not just in memory during test but also study performance.

3.3.2 Relational Memory Errors

Relational memory errors (identity-location misassignments) were evaluated in space and time separately (see **Figure 3.3**), showing that significantly more relational errors were committed in time

than in space (t(171)=15.82, p<0.001). A 2x4 (domain, i.e. space and time, by trial number) repeated measures ANOVA was used to determine if a difference was present across trials. A main effect of domain F(1, 168)=250.62, p<0.001 was found; no significant interaction was present. In space, the number of relational errors decreased significantly between trials 1 and 2 (t(42)=2.75, p<0.01) but not between trials 2 and 3 or 3 and 4 (t(42)=1.86 and 1.36, p=0.07 and 0.18, respectively). In time, there was no significant difference between temporal relational errors between trials 1 and 2 (t(42)=1.3, p=0.20) but a significant difference was present between trials 2 and 4 and 3 and 4 (t(42)=2.70 and 2.32, p=0.010 and 0.025 respectively.
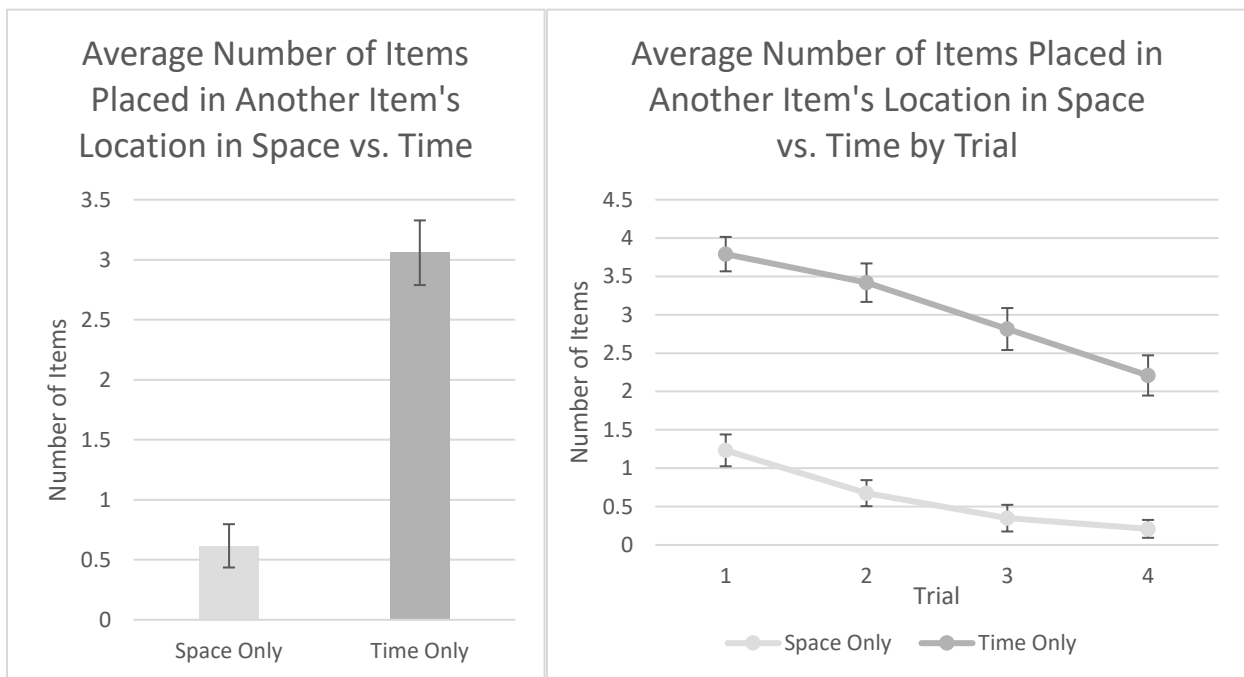

Figure 3.3: Relational memory errors were evaluated in both space and time independently by measuring the number of items which were assigned to another item's location (i.e. a misassignment). The number of these errors in time was greater than the number in space. This difference did not change across trials, though both errors decreased as accuracy improved.

3.3.3 Temporal Context Boundary Effects

To evaluate temporal context boundary effects, only items that had been placed in their correct context were considered. A context boundary effect is computed by normalizing the temporal distance of each item to its within-context and across-context item pair (i.e. the adjacent items in time) relative to the studied distance for each pair (i.e. if the original distance was 4.5 seconds, the distances would be divided by 4.5, thus making a distance of "1" equal to perfect accuracy, while greater than 1 is larger

than normal, and less than 1 is less than normal); then, the normalized within-context pair distance is subtracted from the normalized across-context pair distance to obtain the context boundary effect. This context boundary effect number can either be greater than 0, reflecting a larger distance judgement than was studied, or less than zero, reflecting a smaller distance judgement than was studied. A strong, positive temporal context boundary effect (i.e., stretching of within and expanding of across context distance judgements; see **Figure 3.4**) was observed that when within-context items were placed closer together than across context items in time (paired sample t-test; $t(109)=9.78$, $p<0.001$). This effect was numerically larger for the across context distances, but both within and across differed significantly from 0 ($t(42)=16.02$, $p<0.001$ and $t(42)=10.61$, $p<0.001$, respectively). A 2x4 (within and across by trial number) repeated measures ANOVA was used to compare within and across over trials. A significant main effect of within vs. across was found ($F(1, 106)=80.84$, $p<0.001$) but no interaction effect was present. No significant main effect of trial was present. As no spatial contexts were present, no spatial context boundary effect was assessed.
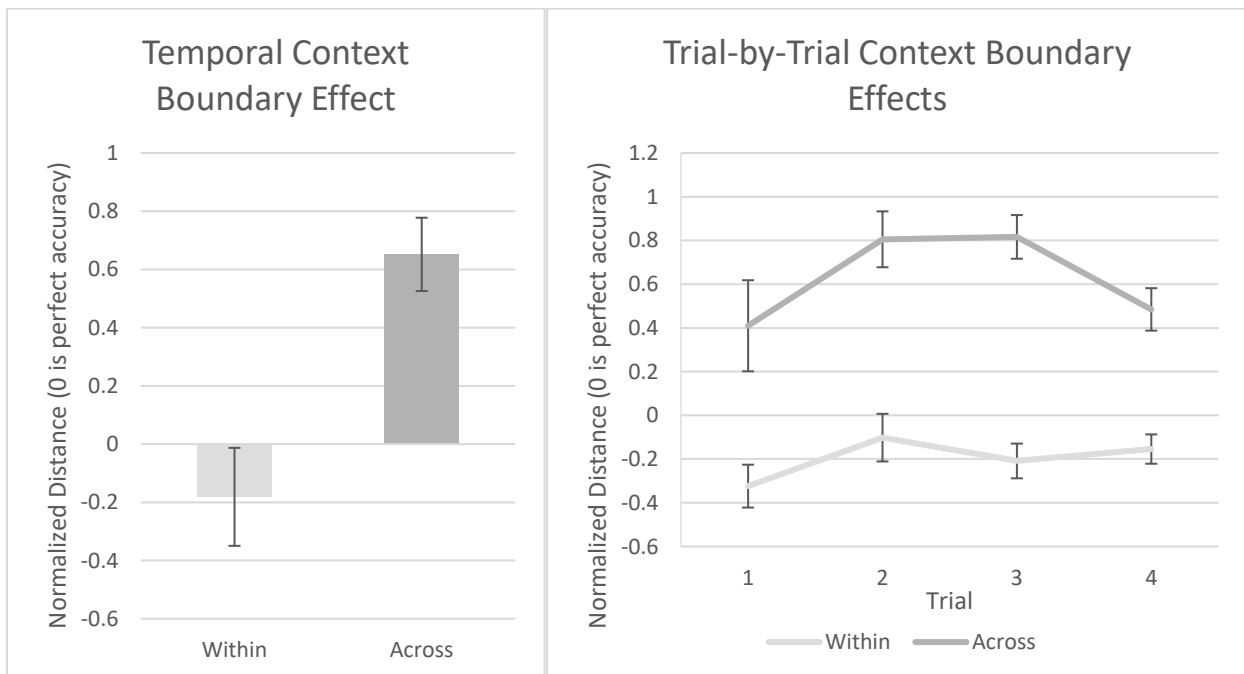

Figure 3.4: Context boundary effects (i.e. the misjudgment of distance of within context items as being closer together than expected and across context items as being further apart than expected) can be seen both in within and across context items. These effects are persistent regardless of the trial number.

3.3.4 Effect of Context on Relational Memory Errors

Finally, when evaluating the identity-location misassignments (i.e. the number of items placed in another item's location) within vs. across contexts, within-context items were misassigned to each other's locations far more often than would be expected by chance ($t(42)=14.63$, $p<0.001$ for across and $t(42)=8.76$, $p<0.001$ for within), suggesting that the contexts biased the relational memory representation, causing confusion between within-context temporal item locations (see **Figure 3.5**). A 2x4 repeated measures ANOVA (within and across by trial number) was performed to examine this effect trial-over-trial. An interaction effect of within and across context misassignment proportion and trial was found ($F(3, 147)=12.12$, $p<0.001$). Follow-up t-tests showed a significant increase in within-context misassignments between trials 1 and 2 ($t(38)=-2.81$, $p<0.01$) and trials 2 and 3 ($t(36)=-4.196$, $p<0.001$) but not trials 3 and 4 ($t(30)=0.53$, $p=0.60$). An opposite pattern can be observed on the across-context misassignment proportion as they are one minus each other. To determine if this effect is present in the absolute number of items being misassigned, another 2x4 repeated measures ANOVA was performed on the number of items to be misassigned within or across context. A main effect of within and across context measures was found ($F(1, 168)=5.74$, $p=0.02$) as well as an interaction between within and across context and trial ($F(3, 168)=11.42$, $p<0.001$). Follow-up t-tests show a significant increase in within-context misassignments between trials 1 and 2 ($t(42)=-2.34$, $p=0.024$) but no significant increase between trials 2 and 3 or 3 and 4 ($t(42)=-1.11$ and 1.50, $p=0.27$ and 0.14, respectively). A significant decrease in across context misassignment was seen between trials 1 and 2 ($t(42)=3.44$, $p=0.001$) and trials 2 and 3 ($t(42)=4.61$, $p<0.001$) but not trials 3 and 4 ($t(42)=1.40$, $p=0.17$.
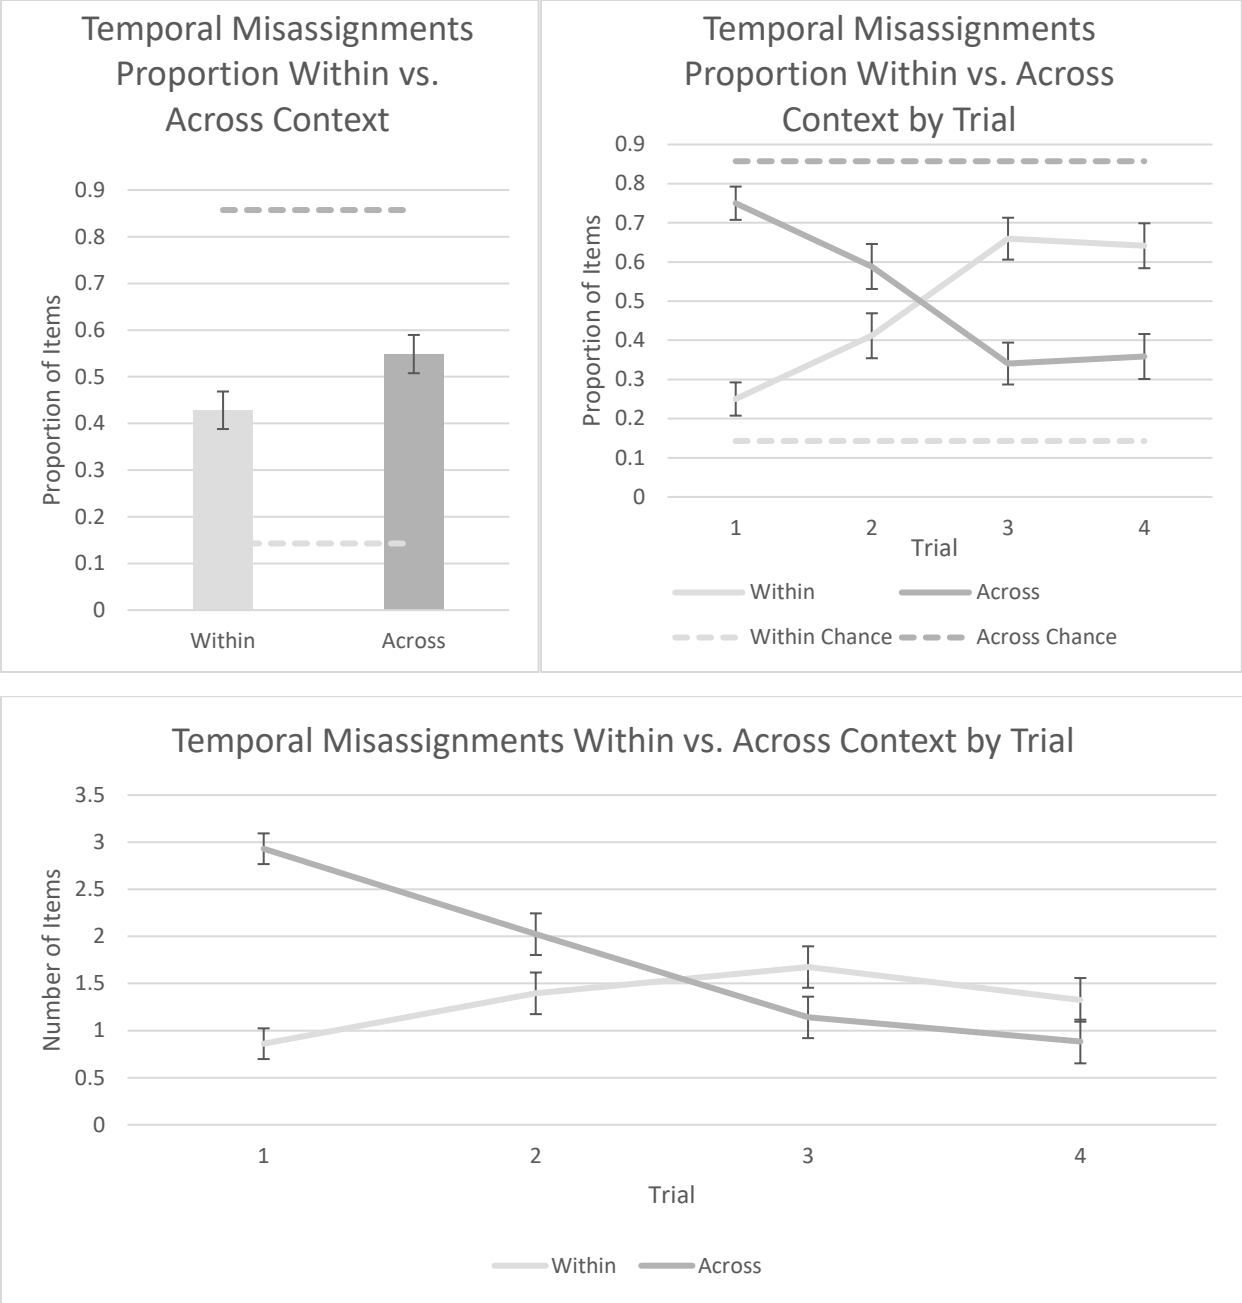
Figure 3.5: The proportion of items which have been assigned to another item's location within vs. across context. Although the within context items should be misassigned to each other's locations by chance at a rate of approximately 14%, we see a substantially higher occurrence of >40%. Interestingly, this effect only gets stronger across trials despite improved accuracy across all non-context based error metrics (including number of misbound items). Note that this same pattern is present when the number of items assigned within versus across are compared directly (rather than as a proportion of the overall errors).

3.4 Discussion

In this work, we used a novel Spatiotemporal Navigation Task to determine if, when space and time are put on a more equal footing, spatial and temporal relational memory errors would be committed with equal frequency. Furthermore, this work examined if, when a timeline of events can be explored bidirectionally, the impact of context on distance judgements in time and the subsequent effect on order memory were consistent with previous findings on context boundary effects. This manipulation is critical for theories of the impact of contextual boundaries on memory as individuals can often re-explore across real-world boundaries arbitrarily. We found that, in this task, relational memory errors in time occurred far more frequently than relational memory errors in space. Context boundary effects were found in time such that within context items were judged to be closer together in time than across context items. Finally, we found that relational memory errors in time were influenced by contextual information as items were misassigned to within context locations far more often than expected by chance. Together, this evidence points to a deep interplay between relational and contextual information in both a spatial and temporal domain, and this evidence suggests that temporal representations in memory may be somehow distinct from spatial representations given the huge disparity between memory in each domain even when they are both able to be explored and sampled similarly. This difference may depend on task demands, and, given that participant performance on temporal measures of relational memory were improving across each trial, it is possible that if more trials were provided, this difference would be eliminated.

The hippocampus is critical for integrating information into relations representations to aid in flexible memory utilization (Cohen & Eichenbaum, 1993). Its role in episodic memory in humans has been established through work with hippocampal damaged patients who have a pervasive inability to create new episodic memories. Additionally, in humans, some evidence shows the flexibility of the hippocampus' potential involvement in real-world representations of space and time, even across long distances (Nielson et al., 2015). In rodents, cells which fire in association with space (i.e. place cells; O'Keefe & Dostrovsky, 1971; O'Keefe & Nadel, 1978), time (i.e., time cells; Kraus et al., 2013), and distance (Kraus et al., 2015) have been identified in the hippocampus and its role in navigation (e.g., Eichenbaum, 2017) has been well established. In line with these observations, some have proposed that space and time make up a scaffolding upon which episodic memories are formed (Ekstrom & Ranganath, 2017). It is not clear, however, if space and time are coded together, as a unified space-time manifold, or if information is bound to each independently of one another. The evidence here cannot completely disentangle these possibilities, however, when using a relational memory error metric which has been

55

shown to be particularly sensitive to hippocampal damage in human patients (Horecka et al. 2017) and the Spatiotemporal Navigation Task we present here, we do find differences in space and time relational memory in healthy individuals such that relational memory errors in time are substantially more common than in space. This suggestions that although these domains may be representable as a unified whole, they are not necessarily so in all circumstances. Task demands may play a large role in determining the degree of interplay between spatially bound and temporally bound information; however, the emphasized domain (time, in this case) may not always benefit from being in focus.

There are several explanations for the phenomena we observe here beyond encoding differences between space and time. First, although both space and time could be explored freely in this task, time remained one dimensional while space had two explorable dimensions. The additional dimension adds more information which must be remembered in order to know spatial locations, but it also provides more information which relational representations may bind. Future tasks may attempt to disentangle these possibilities by restricting space to one dimension or expanding time to two dimensions. Additionally, because time can be explored in any order, it is possible that arbitrarily ordered sampling hurts order memory, increases the chances of misbinding items to other item's temporal locations (i.e. a temporal misassignment error). A similar debate has occurred in the rodent literature in spatial navigation such that distinct hippocampal representations for overlapping routes in a maze are not observed when rats are trained to run in either direction of a Y-maze (Lenck-Santini et al., 2001). This raises an interesting question: can the relative position (i.e. order) of events in time in episodic memory be treated both allocentrically and egocentrically in the same way that spatial position can (see Ekstrom & Isham, 2017 for a discussion of this question)? In spatial navigation, Cognitive Map Theory says that allocentric representations depend on the hippocampus (O'Keefe & Nadel, 1978) with alternate strategies such as view-matching being used to overcome deficits in allocentric memory (Hartley et al., 2004; Lavenex et al., 2014) . However, it is unclear what a view-matching strategy would be when exploring a temporal axis. Place cell firing is related specific to distal cues and is, therefore, seen as strong evidence for this allocentric perspective. Future examinations of temporal navigation should more directly interrogate the potential of allocentric and egocentric strategies in non-spatial domains.

An additional potential explanation for the asymmetry between spatial and temporal relational memory errors is that the temporal axis in this task uniquely contained contextual information (i.e. the spatial axes did not), which may have caused confusion between within-context items (discussed more in future paragraphs). Tasks in which both spatial and temporal contextual information are present might

help to clarify this impact, however, it may not be sufficient to simply include additional types of context to completely disentangle the impacts they have in space and time if spatial and temporal representations are truly distinct.

In addition to the interplay between spatial and temporal memory, this task also allows us to examine the impact of temporal contextual information on memory for a sequence of events. Similar to tasks which exclusively examine order memory, we find context boundary effects such that within context items are placed closer together in time than across context items, even when equidistant. However, this task shows these effects in a continuous timeline of events even when the timeline can be freely explored and the contexts can be restudied freely as often as the participant likes. Moreover, the effect does not diminish when the same short series of events is repeatedly studied. This provides strong evidence that the contextual information is biasing the reconstruction systematically. Critically, of all of the types of errors examined in this task, the context boundary effect is the only error which does not improve with repeated study of the event timeline. Some hierarchical models which attempt to describe how context might be used to organize memory representations suggest that context sits at the top of a hierarchy containing other information such as spatial, valence, and item information (McKenzie et al., 2014). This may explain, in part, why restudy and free exploration don't improve the context boundary effect as the contextual information is being used as a very high level organizing principle for other, more granular information. This may be consistent with high-level theories of the interplay of multiple brain regions involved in contextual and relational learning such as the Complementary Learning Systems model which would suggest that cortical regions which represent and utilize this higher-order information are potentially slower to update than hippocampally-dependent relational information (McClelland et al., 1995). However, recent updates to this view complicate the issue as the speed of acquisition of information in cortical versus hippocampal systems may not be so clear cut (Kumaran et al., 2016). Alternative models which posit a rapid interaction between prefrontal and hippocampal regions in service of decision making (Wang et al., 2015) might better fit behavior in this task as the utilization of the contextual information is necessarily happening in a continuous, real-time manner. However, it is important to note that real-world contextual information is not nearly as clear and discrete as the artificially imposed contexts in this task, and future work might use contextual boundaries (such as night/day boundaries) which are more consistent with real-life experiences.

Finally, one of the most striking results in this work is the interplay between relational information and contextual information. Specifically, we see that the contextual information is related to poorer

performance in order memory even when temporal accuracy improves. This is evaluated by determining the relative proportion of misassignments (i.e., relational memory errors) where an item is placed in another item's location within vs. across a context boundary. If this error were to occur by chance (via randomly ordering the items), we would expect approximately 14% of misassignments to be to a within-context location. However, we find the error occurs more than twice as often as expected within-context. Moreover, within-context misassignment errors become more common (as a proportion of the total number of misassignments) on repeated study both as a proportion of the number of misassignments and in absolute number. This result disagrees with previous accounts which show that context boundaries should impair order memory for across context items and enhance memory within contexts (Davachi & DuBrow, 2015), and shows that, at least when a series of events can be freely explored, context can be related to worse order memory for within context items. There are several explanations beyond the organization of memory due to context which could account for this result. Firstly, it could be that the ability to view a timeline of events in any order results in more confusion for within context items as the viewing order of the context can be different on each exploration. However, it is important to note that participants were not required to view the items in a particular order and could have biased their exploration to avoid this issue. Additionally, if this were the case, we might expect more overall misassignment errors, but this would not completely account for why these errors are so much more common within context than expected unless the context was being used to organize the representation on some level. Another explanation might be that with a limited number of items in the context, not enough information is present to form a reliable relational representation. This explanation may be consistent with the interpretation described earlier that spatial relational information, by having an additional axis within which to bind relations, is easier to reconstruct than temporal relational information. In other words, more information may improve relational memory, to a point. This question could be tested by increasing the number of within-context items to see if this improves within-context order memory.

In summary, with the use of Virtual Reality to create a rich task environment, the importance of contextual information in organizing episodic memory representations can be seen with much more fidelity than previous experiments. Furthermore, contextual information is shown to bias relational memory errors, which were far more common in time than space, such that items were more likely to be misassigned to the temporal location of an item with which it shared contextual information than might be otherwise expected. Spatiotemporal navigation, as presented here, can hopefully serve as a first step in furthering our understanding of the flexibility and structure of relational memory

representations in space and time and clarify the influence of context on different domains of information without the need to treat time differently than we treat space.

Chapter 4: Spatiotemporal Navigation, Sampling, and Information Encoding in Virtual Reality

4.1 Introduction

Gallistel, (1990) said, "Navigation is the process of determining and maintaining a course or trajectory from one place to another." This definition captures a surprising amount of cases which would be commonly considered navigation, and it has been echoed in recent reviews of navigation and hippocampal function in various species (e.g. Eichenbaum, 2017b). Interestingly, however, the word "place" in this context strongly implies that navigation is generally restricted to space, much like "place" is used as the term to describe cells which fire relative to distal cues in space, i.e. "place" cells; O'Keefe & Dostrovsky, 1971). Cognitive Map Theory was introduced by Tolman, (1948) and elaborated upon in the context of hippocampal function by O'Keefe & Nadel, (1978). The theory suggests that, at least in space, that the hippocampus is responsible for maintaining map-like representations, and this theory has been validated to various degrees across many human and rodent lesion studies (Asturet al., 2002; Bartsch et al., 2010; Morris et al., 1982; O'Keefe & Dostrovsky, 1971; O'Keefe, 1991). Tolman himself did not restrict this theory to space; however, later uses of the theory have tended to isolate space in exclusion of any other domains. Despite this isolation of navigation to space, "time" cells in the hippocampus have been identified, which fire relative to particular moments in time (Kraus et al., 2013), and extensive evidence suggests that the hippocampus is involved in mapping far more than just spatial information (Cohen & Eichenbaum, 1993; Eichenbaum et al., 1999; Howard Eichenbaum, 2015, 2017d; Howard et al., 2014; Konkel et al., 2008; Kraus et al., 2015; Ranganath & Hsieh, 2016; Schiller et al., 2015; Warren et al., 2015). It is, therefore, interesting to expand the definition of navigation to include not just movement from place to place, but also from time to time (or, potentially movement between any discrete locations in any domain). The previous chapter introduced a task that can successfully interrogate relational and contextual memory for both time and space, simultaneously, but it did not examine how information sampling (via navigation) during study informed these memory measures. In this work, we demonstrate methods which are applied to the behavioral navigation data collected in the task described in Chapter 3 to examine the relationship between spatiotemporal navigation, learning of temporal sequences, and relational and contextual memory in space and time.

Although temporal navigation is clearly a novel aspect of the spatiotemporal navigation task, spatial navigation is also a large part of the overall task demands. Spatial navigation is a critical capability of all creatures' survival (Wolbers & Hegarty, 2010), and it has been extensively studied phenomenologically in humans (Dudchenko, 2010; Huth, 2013; Wolbers & Hegarty, 2010), animals (Poucet, 1993; Trullier, et

al., 1997), and even robotic navigation (Franz & Mallot, 2000; Levitt & Lawton, 1990) which has, more recently, become an extremely popular topic in machine learning research (Mirowski et al., 2016). Traditionally, in a laboratory setting, tasks such as the Morris Water Maze (MWM; Morris, 1981) have been used to study both rodent and human navigation behaviors (using a virtual Morris Water Maze; vMWM; Moffat & Resnick, 2002), showing that, in older adults, smaller hippocampal volumes relate to longer distance travelled and time taken in the task (Maguire et al., 1999; Moffat, 2009). Recently, more complex behavioral tasks have been attempted, such as observing hippocampal neurons in bats as they navigate in the real world (Geva-Sagiv et al., 2015; Sarel et al., 2017); however, this work is still in its infancy and should certainly be a target for additional investigation. Several models of hippocampal function in navigation have come out of simpler tasks such as the MWM. In rodents, theories of Pattern Separation/Pattern Completion (Rolls & Kesner, 2006; Wilson & Hulme, 1983) have suggested that hippocampal lesioned rodents' deficits in spatial location accuracy are due to an inability to create precise representations of specific locations and differentiate them from other studied locations (i.e. a deficit in pattern separation; Gilbert et al., 1998). Other have expanded on this theory suggesting that after converging "what" and "where" streams reach the dentate gyrus (DG) and CA3 subregions of the hippocampus (Knierim & Neunuebel, 2016), DG specifically performs pattern separation (Leutgeb, Leutgeb et al., 2007; McHugh et al., 2007), orthogonalizing inputs from one another, while CA3 performs pattern completion (Leutgeb et al., 2007; Nakazawa et al., 2002; Neunuebel & Knierim, 2014), auto-associating inputs with previously learned patterns. In humans, the Precision and Binding Model (PBM; Yonelinas, 2013), whose associated evidence shows hippocampal damaged patients are impaired on finding the precise location of a hidden platform but can find the general location (i.e. quadrant; Kolarik et al., 2016; Kolarik et al., 2017), says that the hippocampus is required for high-resolution information binding in spatial navigation. Some models suggest that the long axis of the hippocampus represents progressively more precise information when viewed from anterior to posterior (Fanselow & Dong, 2010; Nadel et al., 2013). In addition to precision of recall, particular strategies of navigation are known to relate to hippocampal activity. View-matching strategies (i.e. those in which a particular perspective on a scene is used as the primary navigation cue) can help overcome deficits in allocentric memory of space via the subject finding an egocentric perspective which allows acquisition of the target (Pecchia & Vallortigara, 2010).

In addition to spatial (and temporal) navigation, this task also critically involves the learning of temporal sequence information. Although temporal sequence memory has been studied extensively and is thought to critically rely on the hippocampus (see Ranganath & Hsieh, 2016 for a review), the ability to

selectively sample and navigate through a temporal sequence and the subsequent impact on memory for the sequence has never been examined. In addition to the hippocampus, other cortical regions such as perirhinal, parahippocampal, and prefrontal cortices may also be critically involved with certain aspects of sequence memory (Hsieh & Ranganath, 2015; Ranganath & Blumenfeld, 2008; Ranganath & Ritchey, 2012a). In particular, it has been hypothesized that the Posterior Medial (PM) network (consisting of parahippocampal cortex, retrosplenial cortex, precuneus, angular gyrus, and ventromedial prefrontal cortex) may play a more pronounced role when the temporal structure is well-learned and meaningful (Ranganath & Ritchey, 2012b). Therefore, it would be reasonable to predict that the PM network would also be involved in exploration of a temporal structure, in addition to the hippocampus. Hippocampal function has also been linked to the ability to use temporal context as an organizational element to allow the automatic retrieval of future elements in a sequence to aid in prediction of future events (Baker et al., 2014; Eichenbaum & Fortin, 2009; Lisman & Redish, 2009; Luft et al., 2015). It is unknown, however, whether or not this predictive functionality will bias behavior when the sequence can be viewed arbitrarily (i.e., with the option to repeatedly review elements on the sequence, maintaining relative order). To assess these behavioral differences, we observe the order of interaction with the items in the temporal sequence to determine if a particular organizational structure is present which might aid in overall learning and memory performance. One final critical finding in the temporal memory literature relevant to temporal sequence learning and retrieval can be found in a sequential retrieval task used by Hsieh et al. where participants were asked to learn the temporal order of objects with either consistent or random temporal positions. Individuals with larger hippocampal sequence representation effects (as measured by voxel pattern similarity analysis) were better able to use sequence knowledge to optimize decision making (Hsieh et al., 2014). Together, this evidence suggests that, similarly to the critical role the hippocampus can play in spatial navigation when arbitrary representations are required, it may play an equally critical role in temporal navigation and subsequent sequence memory under similar conditions. This work attempts to clarify this role through the examination of the relationship between changes in study-time navigation and relational and contextual memory performance.

Several measures of behavior in navigation have been used historically to evaluate the relationship between navigation and memory function, including simple visual inspection of individual paths (Moffat & Resnick, 2002; Sei et al., 1992), heading error (i.e. the sum of the angular distance between the correct heading to the target and the actual heading to the target for a given path; Smith et al., 2013), distance and time to target (Morris et al., 1982), proximity to target (Kolarik et al., 2017; Morris, 1981),

path crossings (Buzsáki, 2005), and even search path classification via clustering (Davis et al., 2010). Each of these measures has advantages and disadvantages worth considering, especially if both spatial and temporal navigation are to be evaluated together. We consider them here as they relate to the Time Travel Task. First, heading error is complicated by the fact that participants could be proceeding to one of many targets in a particular direction. The decided upon event (i.e., the item the participant interacts with next) could be used as a target, but this only makes sense if we assume at least some memory for the event position (not a valid assumption in the first trial for much of the time); therefore, heading error will not be examined in this work. Secondly, path crossing has far less meaning in 3D (2 spatial and 1 temporal axis) than traditional 2D as it is unlikely for a 3D line to cross its self, making it an ineffective measure in this task. Lastly, proximity to target, distance, and time are all potential methods of evaluation, however it is unclear which of these should be preferred for correlation with any particular output measure. Previous investigation into this issue has found that, with a high sampling rate, variable measure of target proximity can be very successful in differentiating hippocampal patient performance from controls (Kolarik et al., 2016, 2017), however, the nature of "proximity" in a spatiotemporal environment is not entirely clear. Future investigations using these measures should first evaluate their efficacy in situations with multiple targets before they are applied to non-spatial navigation tasks.

If none of these traditional measures of path information are appropriate in this case, what measures should be considered? The use of measures of fractal geometry (Mandelbrot, 1967) has shown some success in measuring natural navigation patterns in migration and grazing (Garcia et al., 2005; Gautestad, 2011; Gautestad & Mysterud, 2010). One particular measure of fractal geometry which has shown significant promise is Fractal Dimension (FD), a measure of the complexity of a path (Daugherty et al., 2015). This measure has been related to improved trial-over-trial performance in the vMWM in humans independent of sex and age difference while sex and age tend to be related to more typical measures of navigation such as distance and time taken to perform the task (Daugherty et al., 2015). Moreover, FD has been shown to relate to hippocampal and parahippocampal gyrus volumes in healthy adults when distance and time did not. Finally, the process of computing FD also generates a second component which has not been analyzed in the past in the context of navigation. This second component, known as Lacunarity, can be thought of as a measure of systematicity or compactness of a path. Lacunarity has been applied to several domains involving texture analysis (Plotnick et al., 1996) including cellular morphology evaluation (Smith et al., 1996), lung tumor radiography (Al-kadi & Watson, 2008), and microglia classification (Jelinek et al., 2011). This work intends to determine which, if any, of a subset of measures (enumerated in the Methods section) relate to test-time performance as

measured via the metrics of relational and contextual memory discussed in in Chapter 3. Because of the differentiability of spatial and temporal exploration and reconstruction in this task, we can apply these measures to space and time separately as well as together to determine if there may be differences in spatial and temporal phenomena when they are placed on a more equal footing.

An additional advantage of this task design is the ability to examine learning across several trials of the same complex stimulus environment. Hierarchical Linear Models (HLMs; Gelman & Hill, 2007) are well structured for the analysis of learning in that they allow the comparison of changes in a study variable to changes in a test variable and are used as a primary method of statistical analysis in this work. It may be the case that particular variables relate directly to one another, but this comparison is complicated by two facts: first, on the first trial study, no information about the locations of the items is known by the participant, resulting in a potentially drastically different style of exploration and navigation than on later trials when some memory is present; second, although an overall relationship when collapsing across trials is expected as the trial-over-trial patterns of learning are similar across all study and test metrics, individuals may vary significantly in their relationships between different study and test variables such that a significant relationship is found on, say, trials 2 and 4 but not 1 and 3. The interpretability of these direct relationships is complex at best, and a more straightforward comparison exists in looking at changes in the variables across trials rather than their absolute magnitude.

A final critical feature which will be examined in this task is the impact of contextual information (specifically, contextual boundaries) on subsequent memory. In traditional tasks which have examined so-called Context Boundary Effects (see Chapter 3 for a discussion of these effects), contexts are encountered serially, with no ability to either re-study the sequence and contexts or re-encounter the context boundaries within a given study period (DuBrow & Davachi, 2013; Ezzyat & Davachi, 2014; Polyn et al., 2009). The real world, however, allows ample opportunity for both re-study of contexts and re-encountering of context boundaries, and evidence in event segmentation theory (EST; Zacks et al., 2007) suggests that humans automatically segment continuous streams of information using inferred boundaries (Horner et al., 2016; Sargent et al., 2013). Because this task allows for both re-study of contexts and re-encountering of context boundaries, one critical question addressed in this chapter is whether the number of context boundary crossing has any impact on the context related effects presented in Chapter 3 (namely, the Context Boundary Effect, Within-Context Misassignments, and Across-Context Misassignments).

In summary, this work aims to determine if aspects of spatiotemporal navigation relate to relational and contextual reconstruction memory measures, as well as to assess navigation of the order of events in space and time. This is accomplished by examining navigation data from a Spatiotemporal Navigation Task performed by human participants. Additionally, this work aims to examine how order of navigation of temporal sequences relate to subsequent memory for the sequence when, uniquely, a sequence can be studied in any arbitrary order. Together, this behavioral evidence seeks to establish a foundation upon which future investigation of hippocampal function, navigation, and memory organization can be evaluated together with far greater richness and fidelity than ever before.

4.2 Methods

4.2.1 Participants

43 participants (ages 18-24; mean 20.0 years; 23 males; 13.8 mean years of education) from the University of Illinois Urbana-Champaign participated in the study. In addition to the 43 participants being analyzed, 4 participants were excluded due to not completing the task within the allotted 2.5 hours. All participants signed informed consent prior to participation, approved by the University of Illinois Urbana-Champaign Institutional Review Board and were paid $8 an hour for their time. Note the Navigation data discussed in this Chapter was collected at the same time as the Reconstruction data collected in Chapter 3.

4.2.2 Design and Procedures

This task was identical to the task presented in Chapter 3, so only elements of the task which are critical to Navigation will be reiterated here for clarity.

In this task, participants were placed in a virtual environment measuring 40 meters by 40 meters, bounded by walls on the exterior (see **Figure 4.1** in Chapter 3 for a visualization). The environment resembles a small park with trees, rocks and plants. Participants used the Oculus Rift DK2 and a wireless Xbox controller to interact with the environment. Participants were given a practice to learn the simulation mechanics as well as to gain competency in using the virtual reality hardware. After the practice, 4 study-test trial pairs were given using the same environment as the practice but with new objects and spatiotemporal object-locations. In order to encourage exploration of the environment, an invisibility bubble 10m in diameter was placed around the participant such that any items beyond that distance away from the participant were not visible. An auditory cue was given every time an event happened regardless of whether the item was in view. Item positions were determined via random number generator; however, they were constrained in time to balance for within vs. across context boundary distance. Temporal context information was provided in the form of the simulation boundary

walls changing colors at evenly spaced intervals throughout the simulation. Thus, context boundaries were defined as shifts from one color to the next. The colors yellow, red, green, and blue were used for the contextual information, which changed every 7.5 seconds during the 30 second practice and 15 seconds during the 60 second study and test. Each context had two item events during its duration. The order of the temporal contexts was counterbalanced across participants with half receiving yellow, red, green, and blue and the other half receiving blue, green, red, and yellow. Participants could both walk around the environment and turn their head to look in any desired direction. Temporal navigation was controlled via button press that could reverse the flow of time, at the same speed as the forward flow of time (see **Figure 4.1** for a schematic of the task).

Participants' looking direction and location in space and time were logged at every time point in the simulation (~16ms intervals). If the participant reached the beginning or end of the simulation's time course, time was paused but the participant was still allowed to freely move about the spatial axes in the environment. During practice and study, participants were instructed to find the locations of each item in space and time, specifically when their events happened. To signify they had found the item at its event, there was a 0.5 second window during which participants could see the item's event (if close enough) and press a button on the controller to acknowledge they witnessed the event. Participants were given unlimited time to explore the environment, find each item, and acknowledge they had viewed each event.

During test, participants were asked to reconstruction the spatiotemporal location and identity information associated with each event. See Chapter 3 for more description on test-time data and analyses.

4.2.3 Navigation Analysis Metrics
The study time navigation data are analyzed using several methods, leveraging previous work in spatial navigation but extending these analyses to spatiotemporal navigation and determining if study-time navigation relates to later test performance. Chapter 3 showed the most basic method of analysis, a measurement of the distance travelled in space and time. The analysis of the relationship between these high-level metrics and the various test metrics of interest is presented here, however, there are many measures of exploration and navigation which have been applied to spatial navigation in the past which should also be considered. Specifically, analysis methods of theoretical interest which have been previously linked to hippocampally dependent memory will be examined.
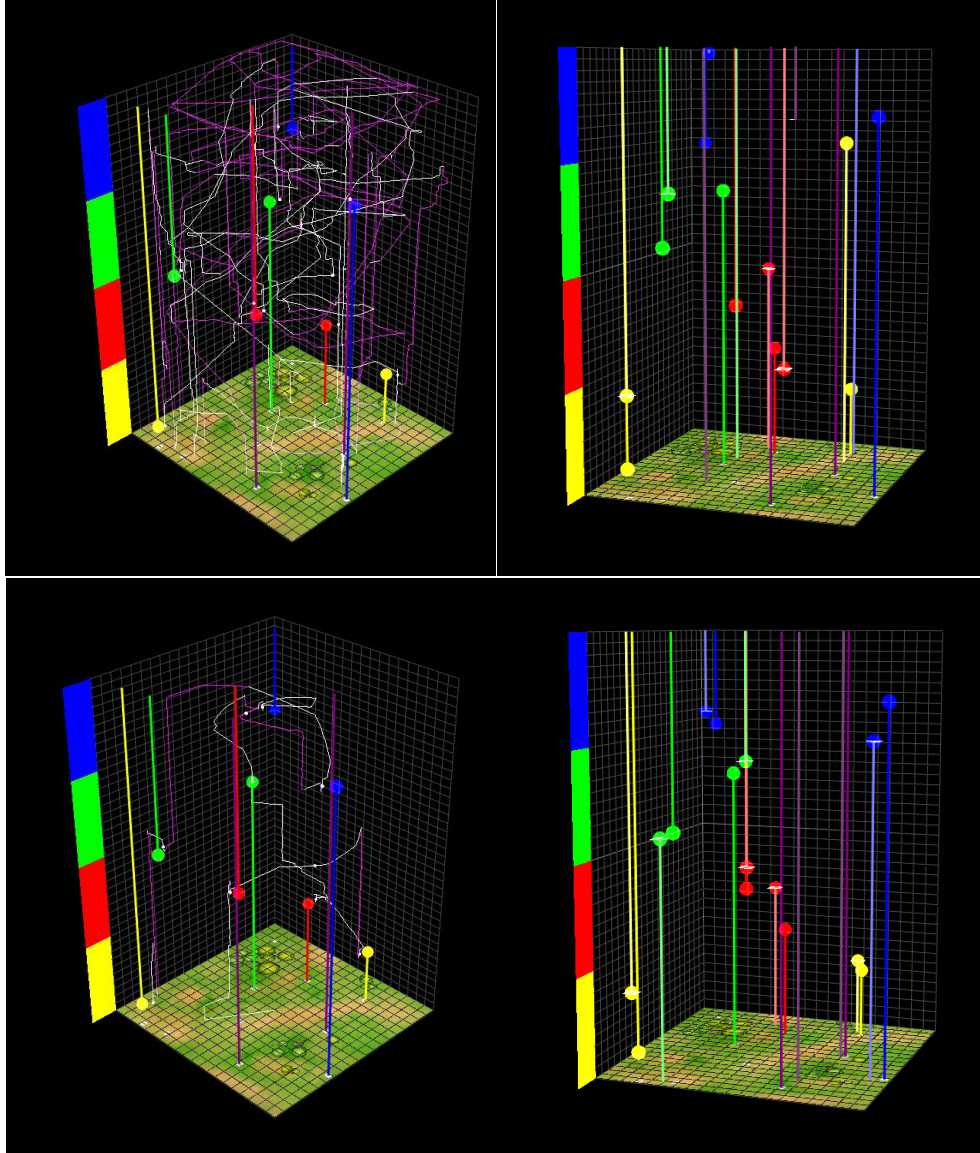
Figure 4.1: A schematic of spatiotemporal navigation and reconstruction results from an example participant's trial 1 (top) and trial 4 (bottom). The 3D volume represents both axes of space (horizontal axes) and the axis of time (vertical axis). The vertical bars on the left indicate the temporal context at each vertical moment in time. The small circles denote the moments in time in which an event occurred and the line denotes if the item associated with the event would have been present at that moment in time. A thin, white navigation line represents forward temporal motion while purple/darker represents backwards. The vertical lines with no dot (purple lines) are objects with no temporal information (always present). The small white dots during study represent moments of interaction with items. The second set of event dots in test represent the reconstructed spatiotemporal locations of each event. Note the path in Trial 1 study is complex and inefficient while the path in trial 4 is substantially more efficient.

Path complexity, as an index of landmark mapping, has been used in ecological research for over a decade, and, more recently a specific measure of path complexity, fractal dimension (FD), has been linked to improved performance in navigation independent of the sex and age differences seen in typical measures such as distance traveled and time taken to perform the task (a virtual Morris Water Maze; Daugherty et al., 2015). Additionally, in the Daugherty study, both the hippocampus and parahippocampal gyrus volumes in healthy adults were related to FD, but not distance and time. As such, FD will be included in the current analysis as a study-time variable that might relate to later test-time variables of relational memory such as misassignments. The FD measure, which up until this point has only been used in reference to spatial path complexity, will be applied to both spatial and temporal path complexity, and will be one of the critical analyses in this work.

In the analysis of FD, the Box Counting Algorithm (see **Figure 4.2**) will be used to determine how changes in scale relate to changes in the measured complexity of the path. Ultimately, this algorithm involves performing a linear fit of the changes in the number of boxes which contain the path as the scale of the boxes change and extracting the slope (aka the Fractal Dimension). In previous studies of paths involving complexity measures, the intercept of this fit (aka Lacunarity) has been discarded, however, this measure is of unique interest as it abstractly represents the degree of systematicity or compactness of the path. A larger Lacunarity is associated with less systematicity, while a smaller one is associated with more. Note that unlike previous implementations of the Box Counting Algorithm, the implementation used here does not require the manual decision of scale parameters (i.e. the range of box scales to be used in the fit). Instead, the scale parameters were picked dynamically (see **Figure 4.2**) by examining a large, log-distributed range of scales and using the standard deviation of the box counting number to remove the asymptotic regions of the measurement (thus, isolating the linear region of interest). In the actual analysis of participant data, the scale range was determined dynamically first, then an identical scale parameter was used for all participants in the final measurement for consistency.
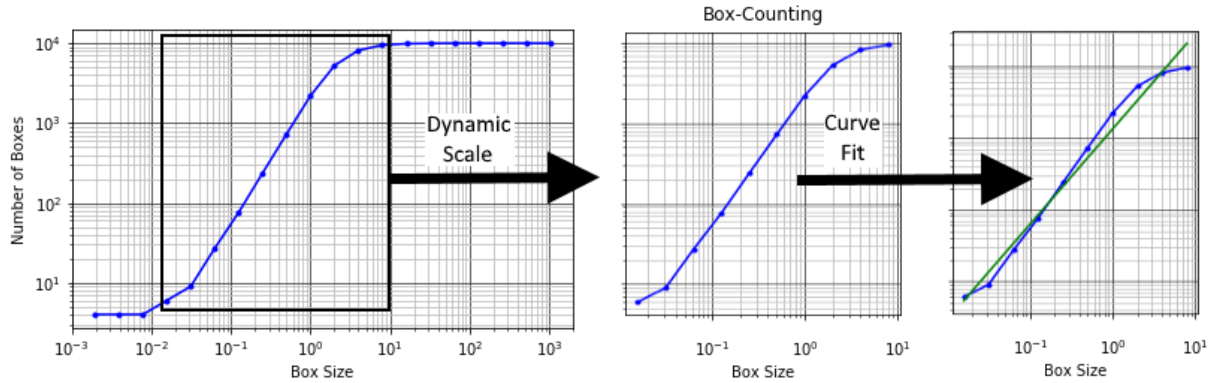
Figure 4.2: Dynamic Computational of Scale Parameters for FD and Lacunarity. In order to avoid issues with determining the precise scale of a path, dynamic scale detection can be used. First, a sufficiently wide range of scales is computed (left). The result should form a sigmoid where the slope at the center represents FD and the intercept of the line fitted to the center is Lacunarity. To dynamically find scale parameters, the Standard Error is computed for the sigmoid and any values exceeding one Standard Error from the Mean are excluded. This shrinks the scale window (center), the linear region which is fit, extracting FD and Lacunarity from slope and intercept of the green line (right).

To see how these two measures (FD and Lacunarity) relate to each other and to the nature of a path, we can observe these measures in the context of filtered random walks. A random walk is one in which a path begins at a particular point (usually (0, 0) in 2D space) and iteratively decides a random direction, moving some distance in that direction. The resulting paths have a fractal-like appearance and are of significant mathematical interest, however by lowpass filtering these paths, we can generate artificial data with various spatial properties in order to examine example paths which have a range of Lacunarity and FD values. **Figure 4.3** shows 4 groups of randomly generated paths of equal length (10,000 points) with various levels of filtering. Each began as a random walk, but were filtered with an order 6 Butterworth Lowpass Filter with increasingly strict frequency cutoff parameters. Notice that for a given vertically-aligned pair of paths, the FD value is very nearly identical, while a large difference in Lacunarity is present (i.e. the path with the larger Lacunarity is less systematic/compact). Similarly, note that horizontally-aligned pairs of paths have similar Lacunarity values but very different FD (i.e. the path with a larger FD is more complex).
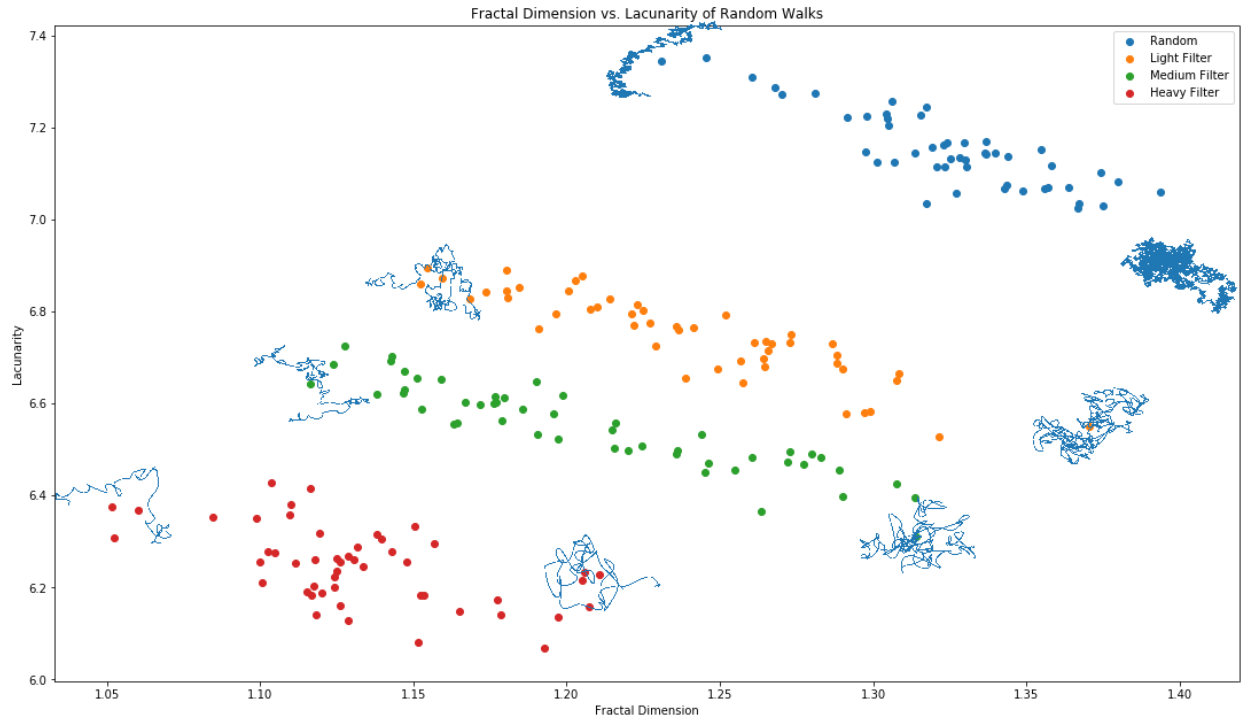
Figure 4.3: Comparison of Lacunarity and Fractal Dimension (FD) of various lowpass-filtered random walks. The blue (top) dots represent fully random walks, while the Light, Medium and Heavy filtered paths were filtered with an Order 6 Butterworth Lowpass Filter with an assumed 30S/s sampling rate and frequency cutoffs of 0.4, 0.2, 0.08. Example paths (blue lines) associated with the extreme points in each group are overlaid on the plot. Although low FD is often correlated with higher Lacunarity, they are not equivalent. Lower Lacunarity within a given filter group shows a more systematic filling of the space while higher FD shows an increase in path complexity. Note that all lines are identical in length.

Beyond measures of path complexity and systematicity, a metric for evaluating the impact of the contextual information on subsequent memory, namely, context boundary crossings, is of particular interest. Most studies of context boundaries do not allow participants to arbitrarily cross boundaries and sample contexts freely (DuBrow & Davachi, 2013, 2016; Ezzyat & Davachi, 2014). As such, it is interesting to determine if additional context boundary crossings are associated with a larger context boundary effect and more relational memory errors due to context.

Table 4.1 provides a summary of the analyses and their associated hypotheses. Note that these are still coarse metrics, collapsing all of the information contained within a spatiotemporal path into a single variable. Additional Order analyses (reviewed in the next section) will be necessary to understand the complex relational information contained within the data.

Measure	Comparison	Support
Spatial Distance	Misplacement and Relational Memory Errors	(Maguire et al., 1999; Moffat, 2009; Moffat & Resnick, 2002)
Temporal Distance	Misplacement and Relational Memory Errors	(B. J. J. Kraus et al., 2015)
Fractal Dimension	Misplacement and Relational Memory Errors	(Daugherty et al., 2015)
Lacunarity	Misplacement and Relational Memory Errors	None (novel metric)
Context Boundary Crossings	Context Boundary Effect and Relational Memory Errors	(Ezzyat & Davachi, 2014; Zacks et al., 2007)

Table 4.1: High-level navigation analyses. Many of these analyses can be run on space and time in isolation of one another as well as together.

4.2.4 Order Analyses

In addition to the simple navigation metrics which evaluate the entire path simultaneously, it may be of particular interest to look at the interactions the participants have with the events (and associated navigation) as related to the order of navigation. Of particular note is if, even when provided with the ability to freely explore time, participants reliably converge towards navigation of the events in the implied "forward" direction of time in a contiguous manner. Previous work on computational models (Howard & Kahana, 2002; Sederberg et al., 2008) and behavioral evidence (Howard et al., 2015; Kahana & Howard, 2005) related to sequence memory would suggest that there is a strong forward bias in contiguity (i.e. a greater probability of remembering the next item in a sequence given an item was remembered), however these models do not allow for the list to be freely sampled (i.e. sampled omnidirectionally and as much as desired). Moreover, this contiguity effect does have a reverse bias of a smaller magnitude than the forward bias (i.e. the previous item in the list is also remembered better than baseline but worse than the next item). It is interesting, therefore, to ask if this bias is still present when participants can freely sample the sequence and if changes in this bias are related to changes in subsequent memory.

Analyzing order to evaluate these biases can be a complex issue. Recency, contiguity, and primacy effects could lead to particular items receiving specialized treatment by chance or due to an overall bias towards certain patterns of reconstruction in time. One potentially robust way to compare reconstruction order is to develop a distance metric for the relative order of events. In other words, given two reconstructed orders. For instance, given A→B→D→C and B→D→A→C, which sequence is closer to the intended A→B→C→D ordering and by how much? Considerable interest has been given to similar problems in comparing the similarity of sequences of DNA, word strings, and compressed data sequences (Birch et al., 2010; Bonham-Carter et al., 2014; Dinu & Ionescu, 2012; Giancarlo et al., 2009). However, many of these involve a much more general case in which no guarantee of uniqueness of

elements in the sequences is provided. Because the task requires each item to be clicked precisely once, instead, the distance metric of interest should focus on three primary properties which might define the differences between two permutations. First, the elementwise-agreement of the lists can be used as a general distance metric. More specifically, the Kendall Tau distance metric is appropriate for this, and it has the added benefit of providing a positive or negative correlation value which, if positive, suggests a closer distance to the forward order and, if negative, suggests a closer distance to the reverse order. Unfortunately, Kendall Tau is incredibly unforgiving to contiguity effects. As such, a second distance metric, known as Contiguity distance, which measures the number of contiguous sublists, can be applied to generate a 2D space of potential distances between a hypothesized and observed order. This distance can be visualized and its probabilities for a fixed length permutation can be calculated explicitly (see **Figure 4.4**). Together these two metrics provide sufficient information to determine if navigation or reconstruction order converge to a particular hypothesized order (specifically, the "forward" order A→B→C→D→E→F→G→H, but also potentially other orders of interest such as H→G→F→E→D→C→B→A or B→A→D→C→F→E→H→G).

In summary, this framework for evaluating order information will be used to observe changes in distance of navigation and reconstruction to the "forward", contiguous order across trials, predicting that both study-time navigation order will trend towards "forward" order across trials (i.e. the low-probability location of (1.0, 0.0) on **Figure 4.4**).

4.2.5 Statistical Analysis

All analyses present in this chapter which compare trial-by-trial performance use repeated measures analysis of variance (rANOVA) using R with Huynh-Feldt correction. For follow-up pairwise comparisons of levels of the rANOVA, Tukey's Method for comparing a family of 4 estimates was applied to the p values. Hierarchical Linear Modeling (HLM; using the nlme package in R) is used in cases of changes in navigation metrics being compared to changes in test performance across trials. In each HLM, normality of the residuals is ensured using a D'Agostino and Pearson's test for normality, and, in cases where required, correction (i.e. log correction) was applied to transform the residuals to normality. Comparisons which collapsed trials together used either paired-sample, two-tailed t-tests or one-sample, two-tailed t-tests (for comparison to expected chance values). All tests used an alpha value of 0.05 to determine significance.
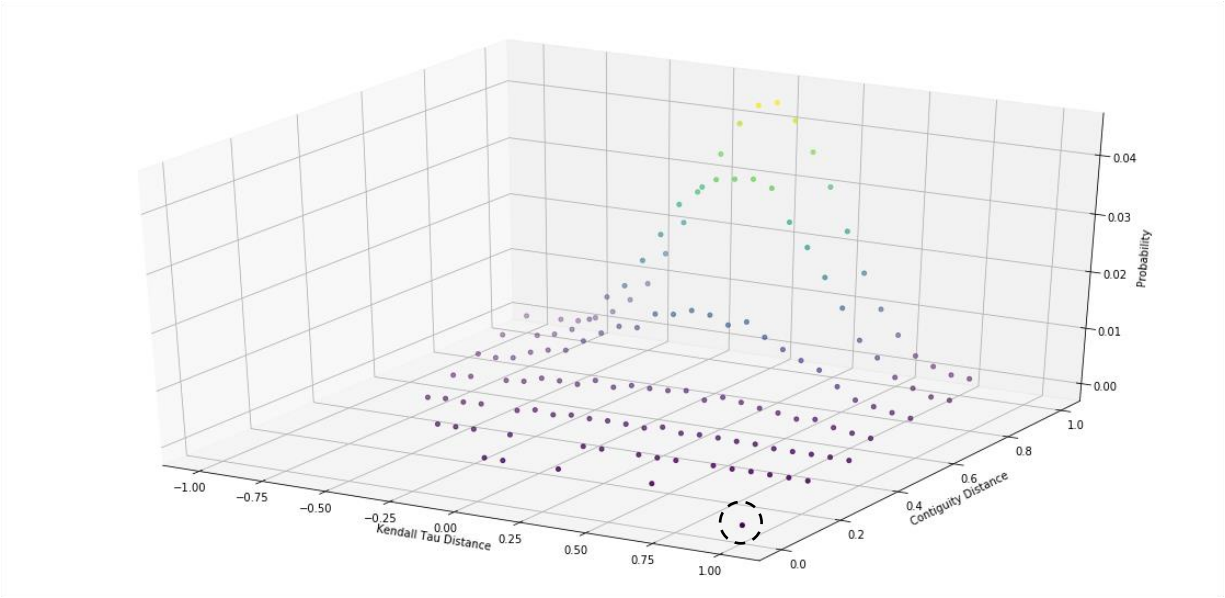
Figure 4.4: Discrete probability distribution for the 2D distance space with 8 items. The x and y axes are the distance metrics (Kendall Tau and Contiguity Distance) to analyze order, and the z (vertical) axis is the probability that that distance would be associated with random permutations. Note that near (0.0, 1.0), very high probability states represent extremely unordered lists (i.e. D→H→B→A→E→G→F→C). On the other hand, four very low probability states exist, with two of primary interest. The point (1.0, 0.0), circled by a black dotted line, represents the order A→B→C→D→E→F→G→H which is the only permutation associated with that point. Similarly, the point (-1.0, 1.0) uniquely represents H→G→F→E→D→C→B→A. The other two unique points which balance the two metrics are ($-\frac{5}{7}, \frac{3}{7}$; G→H→E→F→C→D→A→B; associated with maximum contiguity with minimum element-wise distance) and ($-\frac{1}{7}, \frac{1}{7}$, E→F→G→H→A→B→C→D).

4.3 Results

4.3.1 Changes in Path Navigation Across Trials

In the navigation analysis, 9 key path variables were used to collapse the entire navigation path into a single number. These metrics were: distance travelled in time and space, context boundary crossings, fractal dimension (FD) in time, space, and spacetime, and Lacunarity in time, space, and spacetime. For each variable, a similar shape of learning can be seen (see **Figure 4.5**). Repeated Measures ANOVA reveal a significant change in each variable ($p<0.0001$) for each metric across trials. Follow-up Tukey's Method Contrasts between each pairwise trial combination reveal that for FD space, FD time, FD spacetime, Lacunarity space, and Lacunarity spacetime, all pairwise comparisons were significant (p's<0.05). For time travelled, space travelled, context boundary crossings, and Lacunarity time, all pairwise comparisons were significant (p's<0.05) except between trials 3 and 4.
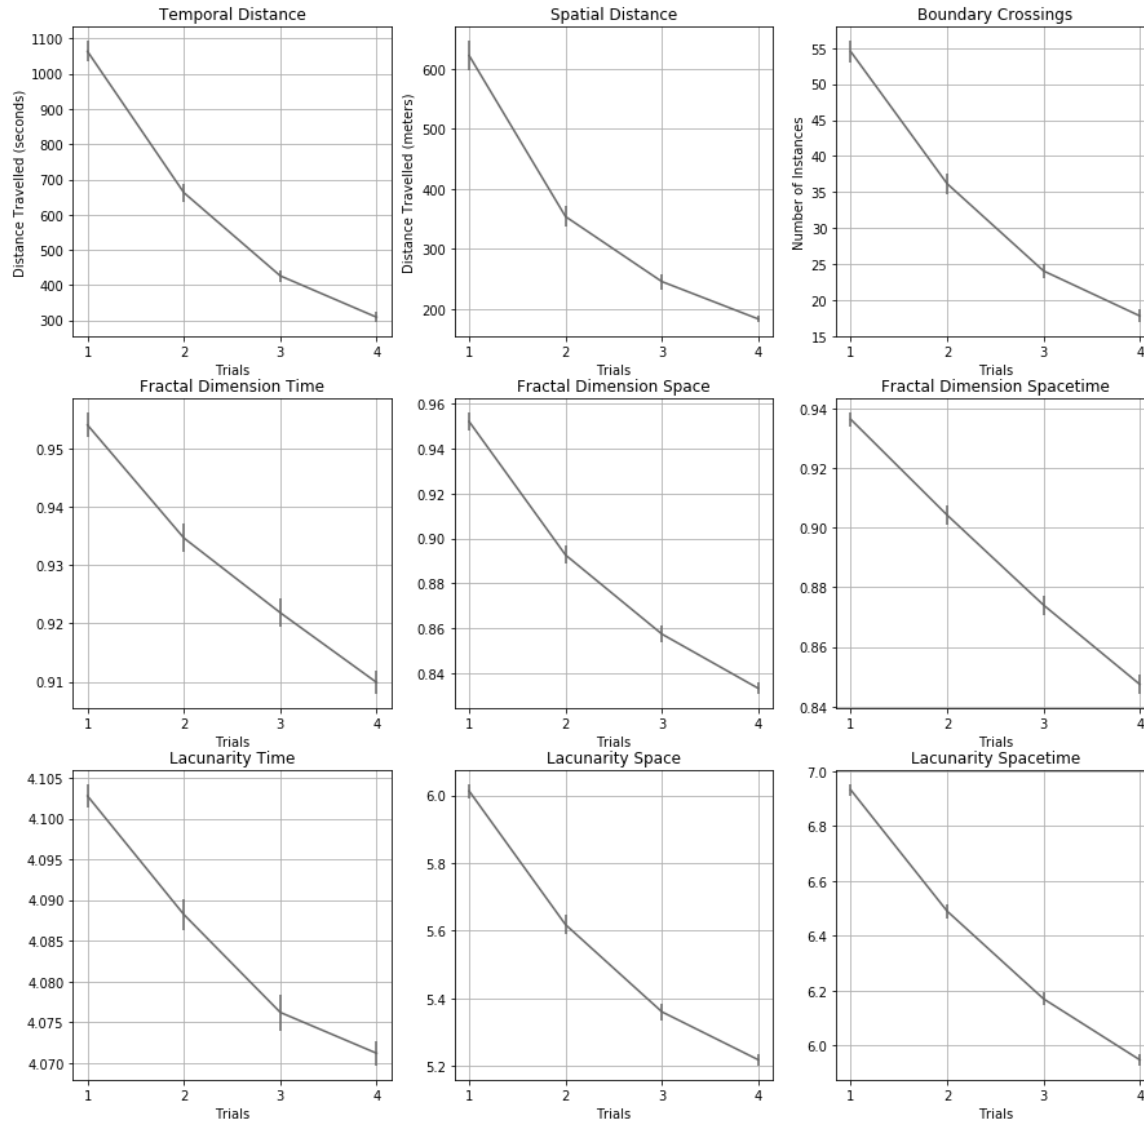
Figure 4.5: Navigation Path Metrics - Note the similarity of learning curves across all path metrics.

4.3.2 Order Analyses

A repeated measure ANOVA was used to analyze the Kendall Tau and Contiguity distances of the order of study of the population across trials (see **Figure 4.6**), and it showed that both distance measures changed significantly across trials ($p<0.0001$). Follow-up Tukey's Method Contrasts between each pairwise trial combination reveal that all pairs at least 2 apart were significantly different from one another (i.e. trials 1 to 3, 2 to 4, and 1 to 4) while all adjacent trials were not significantly different from one another ($p>0.05$) in both metrics. Note that for Kendall Tau, a 0 value is associated with a completely random permutation while values closer to 1 are closer to the "forward" direction (i.e. A→B→C→D→E→F→G→H in the implied temporal order of the simulation). Participants started near

the random point (with a slight positive bias) and trended significantly towards the "forward" point. For contiguity distance, a value of 1 is associated with random permutations while a value of 0 is associated with the "forward" direction. Note that similarly to Kendall Tau, contiguity started slightly below the random point and trended towards the "forward" direction. It is possible that with additional trials, more individuals would have converged to the "forward" direction, but it may also be the case that some other nearby optimal order (or space of orders, perhaps related to the contextual organization of the sequence) would end up being preferred.
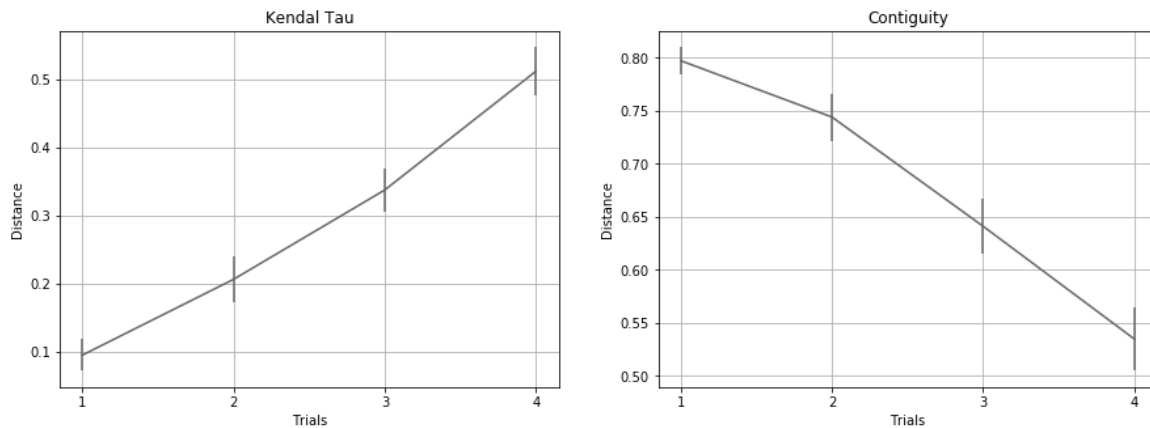

Figure 4.6: Item Interaction Order Metrics - For Kendall Tau, 0 is associated with random order, 1 is "forward" and -1 is "reverse". For contiguity, 1 is associated with random order, 0 is "forward". In both metrics, the trend is towards "forward" interactions with items.

Finally, it is interesting to observe these same findings in the context of the known, prior probabilities of each distance and observe the change in the population within that probability space. **Figure 4.7** shows this comparison where the large blue dot and dotted lines represent the means presented in **Figure 4.6**, while the square markers are individual participants (with the opacity representing the number of participants associated with that point where darker means more participants). The height in this plot represents the probability that, given a random permutation, that particular combination of distance metrics would be present. As can be seen in this figure, not only is the population trending towards the hypothesized "forward" point, but the population is trending away from the high probability region associated with random permutations. Moreover, the proportion of participants who are navigating in precisely for "forward" order is increasing across each trial (up to ~25% by the fourth trial). However, individual differences remain, and some individual's navigation order never deviates substantially from random. This variability in navigation order will be examined as a target for relating measures of study performance to test performance in the next section.

A
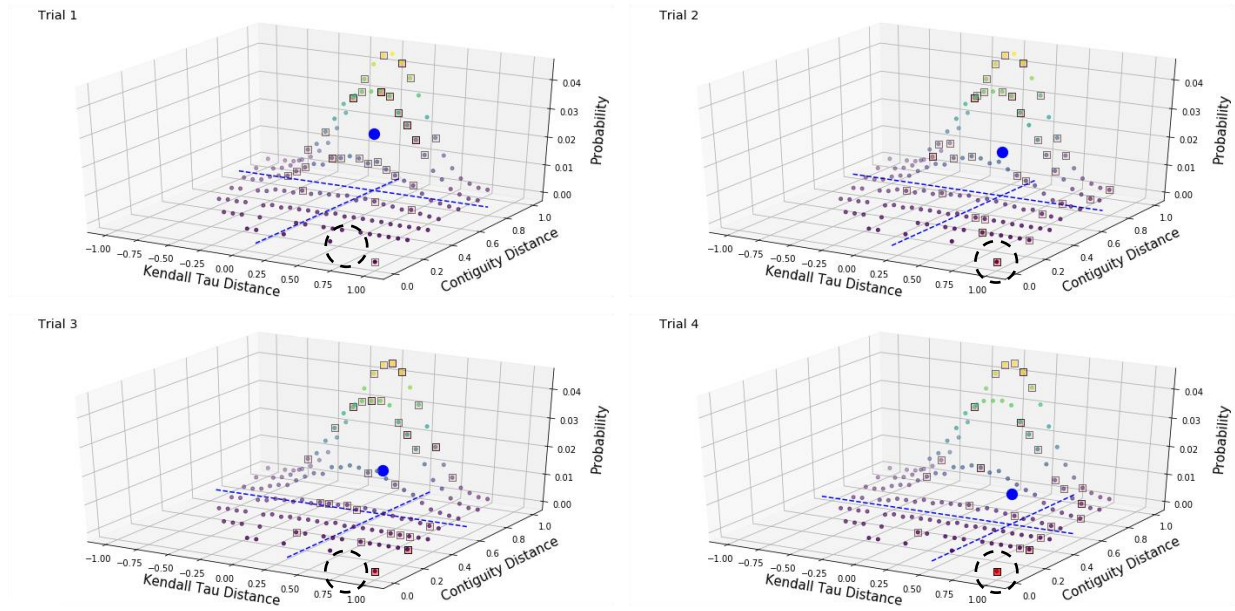
B
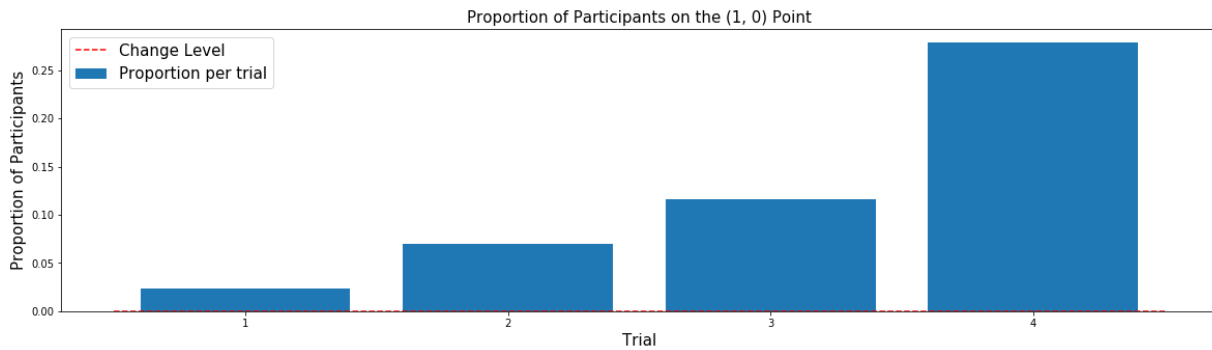
Figure 4.7:

A) Order Metrics - The 4 plots represent performance on the order metrics on each of the 4 trials of the task. The x and y axes represent the distance metrics, and the large/blue dot/blue dotted lines represent the trial means for those metrics (see **Figure 4.6** for just the means). The square markers represent individual participants on each trial. The darker red the square, the more participants were in that particular point. The dotted circles represent the hypothesized "forward" point (corresponding to **B**). Note that despite the group average trending towards the front dot (1.0, 0.0, associated with the "forward" navigation order), there is still significant variability in performance on these measures.

B) The proportion of the participants who were on the "forward" directional point (i.e. A→B→C→D→E→F→G→H) compared to the expected proportion given a random permutation (p=2.5×10^{-5}, corresponding to less than 1 participant given n=43).

4.3.3 Hierarchical Linear Modeling of Study vs. Test Variables

In order to compare the changes in study variables to the changes in test variables, Hierarchical Linear Models are employed. A set of theoretically driven models was chosen which compare variable

76

groupings based on their theoretical relationship to one another. In each case, the random effect component of the model is the trial given the subject, thus changes in a study metric for a particular subject are compared with changes in a test metric for that same subject. **Table 4.2** contains an enumeration of the models under test.

Model(s)	Dependent Variable(s)	Random Effect	Fixed Effect
Navigation vs. Misplacement	Time Misplacement, Space Misplacement	~Trial | Subject	~Time Travelled + Space Travelled + FD Time + FD Space + FD Spacetime + Lacunarity Time + Lacunarity Space + Lacunarity Spacetime
Navigation vs. Relational Memory	Time Misassignment Errors, Space Misassignment Errors	~Trial | Subject	~Time Travelled + Space Travelled + FD Time + FD Space + FD Spacetime + Lacunarity Time + Lacunarity Space + Lacunarity Spacetime
Context Boundary Crossings vs. Contextual Test Metrics	Context Boundary Effect, Within Context Misassignment, Across Context Misassignment	~Trial | Subject	~Context Boundary Crossings
Order Metrics vs. All Test Metrics	All Discussed Test Metrics	~Trial | Subject	~Kendall Tau Distance + Contiguity Distance

Table 4.2: Hierarchical Linear Models under test for comparing changes in study performance to changes in test performance. Note that the Fixed Effects do not include interaction terms, and each model only includes a single dependent variable. Thus, 7 non-order models were tested and 7 order models were tested.

Note that no interactions were tested, as interactions between these variables are not thought to have a meaningful interpretation and in many cases the variables are known to be highly intercorrelated (which is not an issue in Hierarchical Linear Models in general, but it means including interactions can make the model less interpretable with little benefit). Additionally, note that no correction for multiple comparisons was made in these models, as is standard for Hierarchical Linear Models.

4.3.3.1 Navigation vs. Misplacement

When comparing changes in temporal and spatial misplacement in the aforementioned models, in both cases, changes in distance travelled in the domain of interest was the only significantly related variable. Thus, changes in temporal distance travelled across trials are related to changes in temporal misplacement across trials ($p<0.05$, β=0.36, Conditional, Fixed + Random Effect R^2=0.75 and Marginal, Fixed Effect R^2=0.42), and changes in spatial distance travelled across trials are related to changes in spatial misplacement across trials ($p<0.05$, β=-0.42, Conditional, Fixed + Random Effect R^2=0.70 and Marginal, Fixed Effect R^2=0.40). None of the other navigation-related variables (i.e. FD, Lacunarity, or boundary crossings) were significantly related to misplacement in either model. Note that the β value,

which represents how a one standard deviation change in one variable relates to a change in the other variable, has opposite directionality for space and time. This means that while rapid reductions in temporal distance travelled are associated with rapid reductions in temporal misplacement, rapid reductions in spatial distance travelled are associated with slower reductions in spatial misplacement.

4.3.3.2 Navigation vs. Relational Memory

For relational memory errors (namely, spatial and temporal misassignments), the models showed no significant relationship between any study time variable and spatial relational memory errors. One possible reason for the lack of a relationship in the spatial components of test is because of how infrequently these errors occur. However, changes in temporal relational memory errors were significantly related to changes in FD in space (p=0.02, β=-0.46), Lacunarity in space (p=0.008, β=-0.94), and Lacunarity in spacetime (p=0.0003, β=1.91; Conditional, Fixed + Random Effect R^2=0.58 and Marginal, Fixed Effect R^2=0.29). Note that, once again, the β values for the spatial-only components are in the opposite direction from the spatiotemporal component. That is to say, faster reductions in temporal relational memory errors relate to slower reductions in spatial path complexity and systematicity, while faster reductions in temporal relational memory errors related to faster reductions in spatiotemporal systematicity. It is interesting to note that none of the measures of temporal-only navigation related to temporal relational memory. The lack of a relationship may reflect the inability to create expressive or complicated paths in time in the same way as space (either due to the difference in number of dimensions, the requirement that all temporal locations are eventually sampled, or some real difference between the nature of spatial and temporal navigation).

4.3.3.3 Context Boundary Crossings vs. Context Effects

In the context-related models, changes in context boundary crossings significantly related to changes in both the context boundary effect (p=0.02, β=0.24, Conditional, Fixed + Random Effect R^2=0.30 and Marginal, Fixed Effect R^2=0.04), and across-context relational memory errors (p<0.0001, β=0.28, Conditional, Fixed + Random Effect R^2=0.57 and Marginal, Fixed Effect R^2=0.08). However, it is important to note that no significant changes in context boundary effect were seen across trials, which likely explains the low Fixed Effect R^2 for that model. However, this model is far more sensitive to changes in one variable relative to another, suggesting that changes in context boundary effect, when they occur, are significantly related to changes in context boundary crossings. Also note that within-context relational memory errors did not significantly relate to context boundary crossings.

4.3.3.4 Order Metrics vs. Test Metrics

Finally, the order metrics were compared to all test time variables of interest. The significant relationships which were present were exclusive to contiguity. No significant relationships were present between Kendall Tau distance and any test measure. Contiguity significantly related to temporal misplacement (p=0.05, β=0.17, Conditional, Fixed + Random Effect R^2=0.76 and Marginal, Fixed Effect R^2=0.05), temporal relational memory errors (p=0.01, β=0.25, Conditional, Fixed + Random Effect R^2=0.46 and Marginal, Fixed Effect R^2=0.16), across-context relational memory errors (p=0.03, β=0.21, Conditional, Fixed + Random Effect R^2=0.58 and Marginal, Fixed Effect R^2=0.05), and the context boundary effect (p=0.02, β=0.31, Conditional, Fixed + Random Effect R^2=0.26 and Marginal, Fixed Effect R^2=0.12). Note that the directionality of all beta values in this list is the same. Thus, more rapid increases in contiguity across trials are related to more rapid improvements in temporal, relational, and contextual memory measures (specifically context memory errors involving across-context information).

4.3.3.5 Summary of Model Results

For a summary of the results discussed in this section, see **Table 4.3**. Several key findings are worth reiterating concisely. First, changes in navigation metrics at study *did* significantly relate to changes in measures of memory at test. In particular, some separability of the changes in navigation domains (space and time) and changes in memory within those domains can be seen in the misplacement vs. navigation analyses. Also, note that the overall findings are that spatial-only study metrics relate with measures of memory via a different pattern than spatiotemporal study metrics (with opposite directions of β values), consistent with this separability idea. Changes in relational memory errors seemed to relate more to changes in measures of complexity, systematicity, and contiguity of navigation than other metrics. Finally, changes in context boundary crossings were found to be related to both the changes in Context Boundary Effect and changes in across-context relational memory errors.

Dependent (Test) Variable	Independent (Study) Variable	β Value	Condition (Fixed + Random Effect) R^2	Marginal (Fixed) Effect R^2
Temporal Misplacement	Temporal Distance Travelled	0.36	0.75	0.42
Spatial Misplacement	Spatial Distance Travelled	-0.42	0.70	0.40
Temporal Relational Memory Errors	Spatial Fractal Dimension, Spatial Lacunarity, Spatiotemporal Lacunarity	-0.46, -0.94, 1.91	0.58	0.29
Context Boundary Effect	Context Boundary Crossings	0.24	0.30	0.04
Across-Context Relational Memory Errors	Context Boundary Crossings	0.28	0.57	0.08
Temporal Misplacement	Contiguity	0.17	0.76	0.05
Temporal Relational Memory Errors	Contiguity	0.25	0.46	0.16
Context Boundary Effect	Contiguity	0.31	0.26	0.12
Across-Context Relational Memory Errors	Contiguity	0.21	0.58	0.05

Table 4.3: Summary of Hierarchical Linear Modeling Relationships between study and test variables.

4.3.4 Within-Trial Direct Relationships

Finally, linear regression was performed on a per-trial basis on the relationships of interest to determine if any of the relationships found in the HLMs hold directly, i.e. on a single trial between two variables. Interestingly, while most of the aforementioned significant relationships (i.e. space travelled vs. spatial misplacement, time travel vs. temporal misplacement, context boundary crossing vs. context boundary effect, context boundary crossing vs. across context relational memory errors, and temporal relational memory errors vs. spacetime Lacunarity), were directly significantly related on the fourth and final trial (see **Figure 4.8**; $p<=0.01$, $r=0.43$, 0.61, 0.49, 0.38, and 0.40, respectively), very few were significantly correlated on other trials. Across-context relational memory errors and context boundary crossings had a significant relationship on trial 3 ($p=0.01$, $r=0.42$), and time travel and temporal misplacements were significantly related on trials 2 and 3 (see **Figure 4.9**; $p<0.05$, $r=0.42$ and 0.36). All other relationships on other trials were non-significant ($p>0.05$), and, of note, spatial FD and spatial Lacunarity were not significantly related to temporal relational memory errors on any specific trial ($p>0.05$).
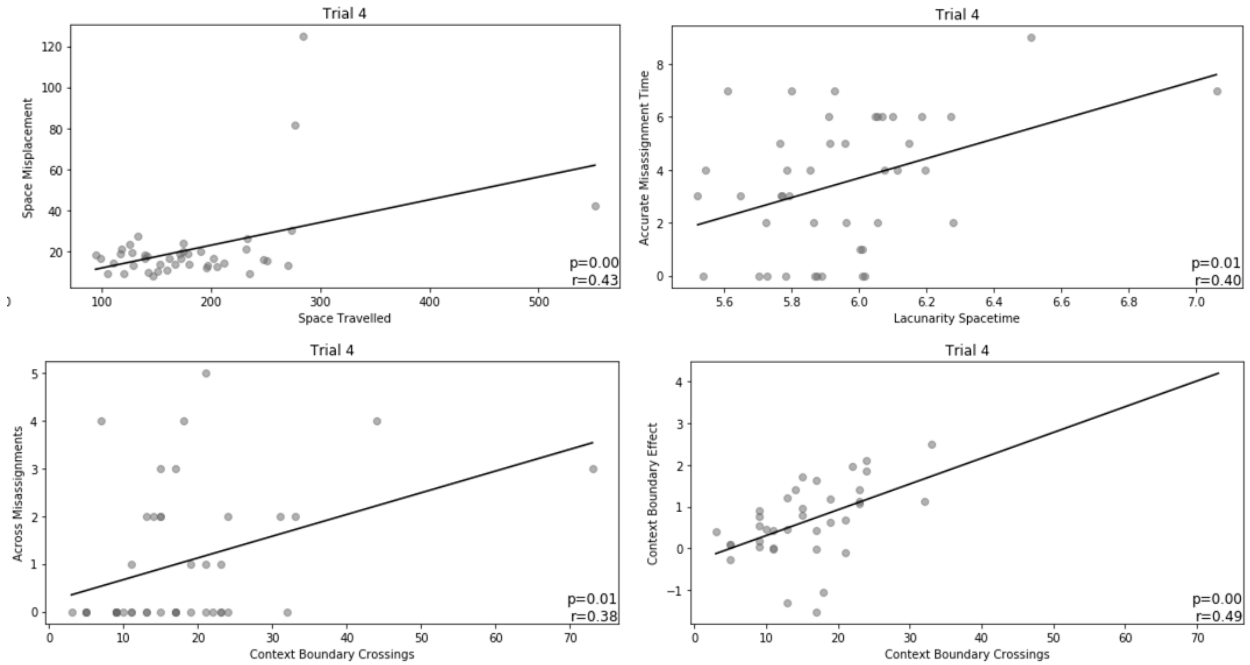
Figure 4.8: Fourth trial relationships between study and test metrics. On the fourth and final trial, several of the relationships which come out in the changes in variables across trials are direct relationships. Here, Space Misplacement vs. Space Travelled, Accurate Misassignment Time (aka Temporal Relational Memory Errors) vs. Lacunarity Spacetime, Across Misassignments (aka Across-Context Relational Memory Errors) vs. Context Boundary Crossings, and Context Boundary Effect vs. Context Boundary Crossings.

Figure 4.9: Time Travelled vs. Time Misplacement. This relationship holds directly on all but the first trial.

4.4 Discussion

Navigation and exploration are critical aspects of the assimilation of information in our day-to-day interactions with our environment. The ability to quickly learn the locations of elements of our environment and flexibly re-encounter those elements at-will underpins a huge number of critical daily activities of both humans and animals. These interactions are often not limited to a static spatial environment, as environments can con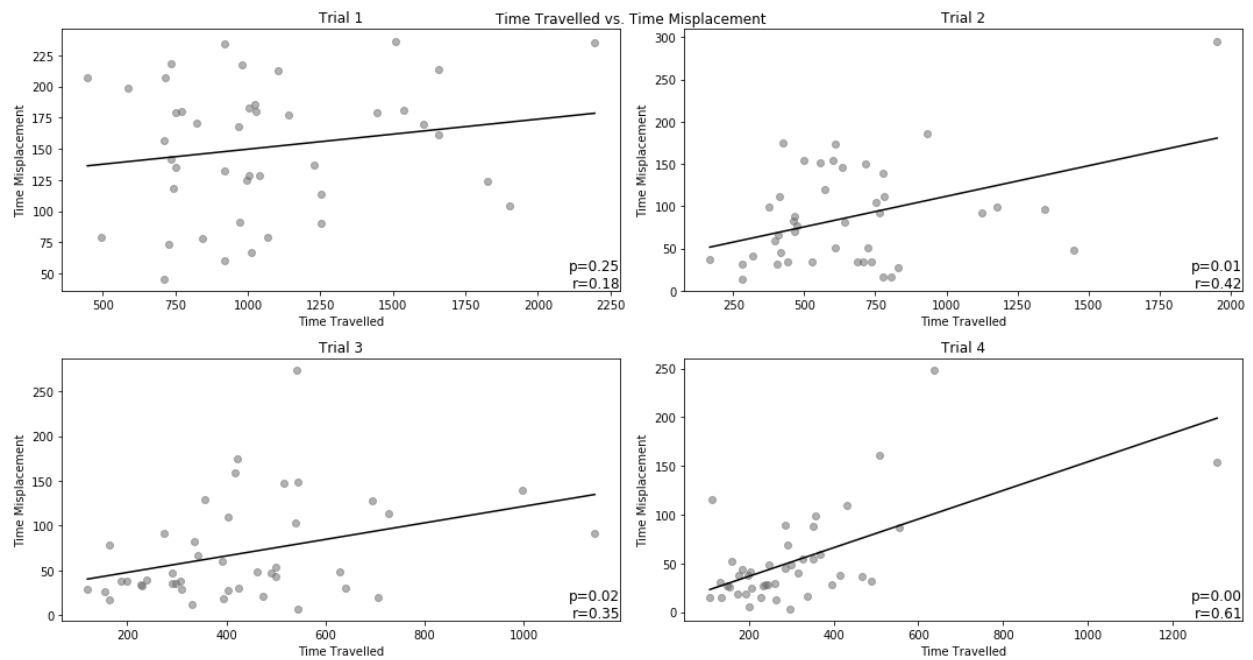tain a complex mixture of spatiotemporal information, hierarchically organized with contextual information. In this task, we examine how changes in navigation of a complex, spatiotemporal environment, organized by context, relate to changes in various aspects of memory (spatial, temporal, relational, and contextual) as evaluated by reconstruction. In this way, several memory phenomena can be studied simultaneously, including the relationship between navigation complexity/systematicity and relational memory, the impact of context boundary crossings on contextual memory biases, and the influence of order of exploration and navigation on relational and contextual memory. In all measures of navigation, and most measures of reconstruction, the participants in this task significantly improved across trials. It is the nature of this improvement and the relationships between study and test improvements that are of special interest, as these relationships reveal interesting new details on how navigation and reconstruction relate in services of the organization of memory in general.

Several critical relationships were observed in this chapter. First, the simplest relationships are those between changes in the overall navigation distance in each domain (space and time) and the changes in the overall misplacement (Euclidean reconstruction distance) in the respective domains. Note that in these models, all navigation variables of interest (i.e. FD in space, time, and spacetime, Lacunarity in space, time, and spacetime, and distance in both domains) were included, so it is fascinating to note that only the particular domain in study was related to the corresponding domain in test at the coarsest test analysis level. This suggests that improvements in navigation (or lack thereof) within each domain might be in some way independent such that it is possible to improve in spatial misplacement by focusing on spatial navigation optimization while neglecting time or vice versa at a coarse level. This evidence of some manner of separability or independence is made more compelling by the fact that the directionality of the relationship is opposite for space (slower changes in study behavior related to faster changes in test performance) as it is for time (faster changes in study behavior related to faster changes in test performance), and that this directionality is consistent across the all analyses of space and time which were statistically significant. This aligns with the evidence from Chapter 3 showing that measures of spatial and temporal relational memory (i.e. spatial and temporal misassignments) also had

significantly different magnitudes, suggesting some manner of separable representations of space or time in this task. Although alternative explanations exist (such as differences in the perceptibility of time vs. space, and the relative familiarity with spatial versus temporal navigation), these results converge with other findings (Dutta & Nairne, 1993; McKenzie et al., 2014; Rondina et al., 2016) that space and time my, under some circumstances, be represented separably.

No relationships were found between any changes in navigation variables with spatial relational memory (likely due to the relative infrequency with which these errors occur in this task), but significant relationships were found between changes in spatial FD, spatial Lacunarity, and spacetime Lacunarity and changes in temporal relational memory, suggesting that changes in systematicity and complexity of exploration of the environment critically relate to temporal relational memory learning (i.e. changes in temporal relational memory accuracy on restudy). It is worth noting that there were differences in the directionality (as seen by the β values in **Table 4.3**) between spatial-only components and spatiotemporal components with temporal relational memory. In both the simple navigation/misplacement case and the relational memory case, changes in spatial-only components negatively relate to changes in misplacement and relational memory errors. Rephrased, this means that a more rapid change in spatial navigation complexity and systematicity relate to a slower change in misplacement and relational memory. However, critically, spatiotemporal Lacunarity (i.e. systematicity of spatiotemporal navigation) has a large, positive β value, suggesting that rapid changes in spatiotemporal systematicity relate to rapid improvements in temporal relational memory.

In addition to spatial and temporal misplacement and temporal relational memory accuracy, the relationship between changes in context boundary crossings and changes in context-related memory phenomena can be directly examined in this task. The ability to reexplore temporal contexts is a unique contribution of this task, so examination of the relationship between interactions with context boundaries and subsequent test phenomena is of particular interest. Moreover, the ability to restudy the same context-dependent stimuli repeatedly allows examination of how changes in these interactions relate to changes in test performance. In these HLMs, we observe that changes in boundary crossings significantly relate to changes in the context boundary effect, as well as changes in across-context relational memory errors. In other words, participants whose instances of boundary crossing decreases rapidly should also show a rapid decrease in context-related memory biases and errors. Note that this relationship is unique to the across-context relational memory errors, which may potentially be due to the unique way in which within-context errors actually increase in this task across trials. This

would suggest that the boundaries are acting more as a discriminatory influence on the relational and contextual memory in this task rather than directly aiding in (or harming) the within-context relational representations. This is consistent with the accounts of event segmentation theory (Sargent et al., 2013; Zacks et al., 2007) and computational models of contextual shift (Howard et al., 2005; Polyn et al., 2009), which posit that event boundaries act to segment sequence representations. However, this finding is not entirely consistent with evidence suggesting event boundaries may improve memory for within-context sequence representations (Heusser et al., 2018). Future studies using the Spatiotemporal Navigation Task may seek to resolve this contradiction by including more items within each context, performing a binary decision test (instead of reconstruction) using a similar free-navigation paradigm for study, or increasing the numbers of contexts being sampled to determine if these factors are critical for improving within-context sequence memory.

Finally, several HLMs were used to explore the relationship between order of study and all relevant test variables from Chapter 3. Changes in contiguity of interaction with items in time uniquely related to the changes in various temporal and contextual test variables (i.e. temporal relational memory errors, temporal misplacement, the context boundary effect, and across-context relational memory errors). In other words, participants who showed more rapid adoption of a forward-ordered study pattern showed a more rapid reduction in errors of all non-spatial types (with the except on within-context relational memory errors). Although many explanations of this relationship are possible, contiguity, in this case, may have been aiding in forming a more map-like, topological representation of the temporal and contextual structure of the environment. Previous work examining sequence retrieval has demonstrated that participants with larger hippocampal sequence representation effects were better able to use sequence knowledge to optimize decisions (Hsieh et al., 2014). It is possible a similar optimization (or simplification) of environment interaction is at play in this task such that participants use the inherent structure of the simulation (including the temporal contexts) to create an organized representation of spatiotemporal information. We know that temporal context can be a critical organizational element in sequence memory (Eichenbaum, 2013; Howard et al., 2005; Howard & Kahana, 2002; Howard & Natu, 2005) and is also dependent upon the hippocampus (Davachi & DuBrow, 2015; DuBrow & Davachi, 2016), therefore it is especially interesting that changes in the contiguity of navigation order relate to both the relational and contextual memory measures in this task. Although only a quarter of the participants ultimately explore in precisely the forward order, many more participants ended up in highly contiguous states, while a non-negligible proportion of the population never deviated from a

pseudo-random order. This individual variability in the use of task-specific temporal structure is a potentially interesting target for future investigation.

It is interesting to note that, although previous paragraphs discussed the complexity of interpreting individual trial or direct relationships between study and test variables, many of the relationships discussed in prior paragraphs show significant direct correlations for the fourth, and only the fourth, trial of this experiment. This may reflect the fact that for many of these measures, once sufficient opportunities have been provided to participants to learn the reconstruction, the individual variability along the critical axis (and not on the orthogonal axis) reaches a point where the direct effects can be disentangled. That is, some individuals may not have been as successful in learning the reconstruction by the final trial, and know it is their last opportunity to attempt an accurate reconstruction, their final trial navigation represents a last-ditch attempt to assimilate the required information. While participants who learned the environment very successfully by the fourth trial need no such additional sampling and navigation attempts. If this were true, it would be expected that these effects might go away had more trials been given and the lower performers had a chance to catch up the higher performers. Of course, this is just one of a myriad of possible explanations, and future examinations of learning in similar paradigms will be necessary in order to determine the reason for these effects on the fourth trial.

The current work seeks to extend and elaborate upon examinations of navigation by other researchers in a variety of ways, leveraging the rich, Spatiotemporal Navigation Task discussed throughout this and the previous chapter. Insights from this task relate to a huge variety of theoretical models of memory, including relational memory theory (Cohen & Eichenbaum, 1993), cognitive map theory (O'Keefe & Nadel, 1978), precision and binding model (Yonelinas, 2013), models of temporal context (Howard & Natu, 2005), and models of extrahippocampal contributions to navigation (Ekstrom et al., 2017), to name a few. Before examining the specific findings of this Chapter in the context of these models, it will be useful to consider perspectives on cell-level representations of information coding and how they might ultimately relate to a task in which navigation is performed in a domain other than space. Although it is widely accepted that place cell firing is related to allocentric information relative to distal cues in the environment, other cell types which fire related to a variety of other relations have been identified. These include grid cells which fire in accordance with periodic spatial location arrays, cells that fire in accordance with goal vectors (Sarel et al., 2017), head-direction cells (Ranck Jr, 1984; Taube et al., 1990), time cells (Kraus et al., 2013), distance cells (Kraus et al., 2015), time-distance cells (that

fire in relation to both variables), and cells that signal borders and speed (McNaughton et al., 1983; Hartley et al., 2013; Kropff et al., 2015). Some have theorized other, more generalizable cell models such as "context cells," which are proposed to exist in the recurrent collaterals of CA3 cells and activate at different points in a sequence, underpinning sequence learning (Hasselmo & Eichenbaum, 2005; Levy, 1996). A naïve view of these cells might say that various cell assemblies in the hippocampus are dynamically mapping to basic signals of navigation (and only navigation). However, substantial evidence links hippocampal function to far more than just navigation (see Cohen & Eichenbaum, 1991 for a review of these arguments). A much more robust perspective would suggest that the hippocampus performs arbitrary binding and does so with any and all relevant domains of information (or perhaps, as suggested by some, preferring space and time but utilizing task-relevant domains; Ekstrom & Ranganath, 2017). In this view, the hippocampus is considered one component in a much wider "navigation network" (perhaps centered on either the retrosplenial cortex or posterior parietal cortex depending on the degree of allocentric or egocentric information being utilized during the task; Ekstrom et al., 2017) while also being a region that critically supports non-navigatory processes (Coronel et al., 2012; Eichenbaum, 2017d; Konkel et al., 2008; Kumaran et al., 2016; Tavares et al., 2015; Warren et al., 2015). Electrophysiological evidence in rodents supports this as hippocampal networks have been shown to associated both spatial and nonspatial dimensions (McKenzie et al., 2014). Indeed, in rodents, humans, and monkeys, recording studies have shown hippocampal involvement in mapping the time course of events independent of space (Kraus et al., 2015; Kraus et al., 2013; Naya & Suzuki, 2011; Pastalkova et al., 2008; Paz et al., 2010; Spiers et al., 2001). Perhaps, then, in the spirit of "context cells", a more generalized definition for many of these information-selective cell types might to call them "relational cells", and then clarify what relation(s) are of specific interest to that cell at that moment. Within this framework, it then becomes critical to examine complex processes like navigation in a wider context, necessitating the use of a greater variety of metrics of performance in navigation. This work seeks to do just that, as well as to understand the relationship between these different measures of navigation and aspects of memory which are known to be hippocampally dependent.

Unfortunately, coarse navigation metrics such as those in **Table 4.1** miss an enormous amount of the fidelity in this task; however, much of the moment-to-moment decision making in this task is currently beyond the reach of our ability to make inferences. Even in the simplest cases (such as the order of study of events), the inference of the intention of the participant becomes a critical one. This work, however, will likely inform future studies and highlight two primary strategies for future investigation: First, careful task design can strengthen the power of the analyses by restricting some possible

strategies and intentions. For instance, restricting the spatial axis to 1D might help us determine if the differences in spatial and temporal memory are due to the spatial domain containing more information or due to the additional manipulations in the temporal domain. Secondly, as artificial intelligence via Deep Learning continues to advance, we are presented with the unique opportunity to use these systems as points of comparison to real neural systems. Deep Learning systems (such as Asynchronous Advantage Actor-Critic, A3C models; Juliani, 2016; Mirowski et al., 2017) can now perform astounding navigation tasks such as solving mazes using only visual input, motor output, and a simple reward signal (i.e. with no special knowledge of the environment or system beyond some sense of "good" and "bad" results of behavior). These systems may be modelled after biological systems, but they operate in very distinct ways. Nonetheless, they are now capable of performing many of the same complex navigation tasks that simple organisms can, and moreover, we can manipulate them in any way imaginable (and as often as we would like). This provides a fantastic test platform for theoretical approaches to navigation analysis which should not be ignored by cognitive science. Collaborations with deep learning researchers and cognitive scientists do exist (viz., Kumaran et al., 2016), but they are not nearly as common as one would hope given the power of these systems to make and test predictions about information processing and behavior.

In summary, we believe this task and the associated methods shed light on the complex interactions of navigation and memory organization in multiple domains and in the presences of contextual information. Individual improvements in spatial and temporal navigation relate to improvements in memory in those domains separably, suggesting that spatial and temporal representations may in some way be separable in this task, at least at the coarsest level of test analysis. Temporal relational memory improvements, however, are uniquely tied to changes in navigation complexity and systematicity, suggesting a critical and complex interplay between *in-the-moment*, *memory-guided decision making* (potentially critically tied to hippocampal representations) and subsequent relational memory efficacy. Further, context boundaries may act as more of a discriminatory feature (at least in this task) than one used to strengthen within-context relational memory organization accuracy, even when allowed to be re-experienced freely. Lastly, a preference towards exploring an otherwise temporally-flexible environment in the implied, forward order with increasing contiguity may be a critical strategy in improving temporal, relational, and contextual memory organization. Taken together, this task establishes both an experimental and analytical framework upon which future investigations (such as imaging or patient work) may more directly probe the relationships between behavioral correlates of navigation and memory in specific brain regions, such as the hippocampus, prefrontal cortex, and

retrosplenial cortex. The evidence presented here show that changes in relational and contextual memory performance are critically tied to changes in navigation behavior, and biases due to different levels of systematicity, complexity, contiguity, and context boundary interactions all contribute to subsequent memory. Moreover, this novel framework can be applied to domains of information beyond just spatial navigation to elucidate the relationships between memory and navigation in many other domains.

Chapter 5: General Discussion

5.1 Summary of Results

The previous three chapters have shown a variety of data from both hippocampal damaged individuals as well as healthy adults, which together establish a methodological framework which can be used to analyze and understand behavior in reconstruction tasks in any domain. Critically, certain error metrics associated with the ability to retain arbitrary relational information were identified as being related to hippocampal function via observing differences in these metrics between hippocampal damaged patients and controls. These same errors were then evaluated in healthy individuals in other domains to determine if two, omnipresent domains, i.e. space and time, have any differences in the relative commonality of these errors, suggesting representational differences between these domains. In order to study space and time simultaneously and on a more equal footing, we allow sampling of both spatial *and* temporal information, via "time travel", thus balancing the ability to control what spatial and temporal information to sample. Finally, because this manipulation provides a new axis upon which sampling can occur, sampling behavior can also be studied independently in each domain. Taken together, the methods and results here provide a foundational framework for studying memory organization and navigation across more domains than just space (i.e. temporal spaces, perceptual spaces, social spaces, and auditory spaces) and with contributions from types of information beyond spatial information (i.e. contextual and sequence information). In addition to finding that, within space, only certain *types* of relations are impaired in hippocampal damage, these same relational *types* can be found to show differences in spatial versus temporal domains and relative to contextual boundaries, suggesting representational differences between these domains in the hippocampus as well as a critical role for contextual information in the organization of these domains. Finally, changes in navigation behavior in both time and space relate to changes in measures of temporal relational and contextual memory, suggesting that hippocampally dependent relational memory may be critically related to in-the-moment decision making in navigation when the task involves learning of arbitrary event information in multiple domains. All-together, we believe these results show that the role of the hippocampus in navigation and memory is in the binding of *arbitrary* information, specifically, regardless of the domain from which that information is derived, and that this binding is performed in service of both long-term, episodic memory and online processing.

In order to examine these theories of hippocampal function, in the second chapter of this document, a framework for the analysis of reconstruction errors was introduced which attempted to assess many of the same informational metrics of previous work (Jeneson et al., 2010; Watson et al., 2013) but from a

first-principled approach with minimal assumptions. Within this framework, misplacement, i.e. the sum of the Euclidean error in space in reconstruction, is treated as a coarse but fundamental measure of error. Misplacement, under this assumption, can then be broken into different error metrics in accordance with aspects of information which were not present in the reconstruction. Most importantly, identity information (i.e. the information which allows items to be distinguished from one another) is removed from the reconstruction during the first step of analysis to determine aspects of the spatial relationships between items and environment which are impaired. This includes evaluating the overall amount of translation, scaling, and rotation shared by items as well as the remaining local misplacement and accuracy for each item's location information. Once these spatial errors are determined, the identity-related errors can be reinstated within this informational framework to determine if item identity is being misrepresented in a systematic way. The key finding in this chapter showed that hippocampal damage impairs the ability to bind arbitrary identity information to locations, but it does not impair the ability to remember the locations themselves. Moreover, this impairment in identity information is fundamental, in that there is no evidence that it is due to compound identity errors in which multiple identities are systematically swapped/cycled amongst different positions. Although swaps and cycles do uniquely occur at low item set sizes in hippocampal damaged patients, once both patients and comparison participants began making the prerequisite error (i.e. any identity-location misassignment), there was no difference in the number of compound errors between groups. This finding is distinct from previous work on spatial reconstruction which showed significant differences in swapping at all set sizes (though most significantly at low set sizes; Watson et al., 2013). The reason for this difference is in the increased sensitivity and specificity of the methods presented in this second chapter over those previous methods.

In the third chapter, the methodology established in the second chapter was applied to a more complex task in which both space and time were navigable domains, simultaneously; this task is known as a Spatiotemporal Navigation Task. In this task, participants were asked to reconstruct events in space and time while being allowed to explore both domains, simultaneously. Although space was explored via joystick, time was also allowed to be explored via a "time travel" button, which allowed the reversal of the direction of the flow of time. Participants performed this task four times on the same timeline of events, and their performance on spatial, temporal, and item identity information improved across all trials. The identity-location misassignment metric introduced in the second chapter was used to assess difference between spatial and temporal relational memory across trials and found that significantly more temporal misassignments (i.e. events placed at the moment in time of another event) were

committed that spatial misassignments. Additionally, the task contained 4 distinct temporal contexts (denoted by a change in the background color of the environment every 15 seconds) which could be used to help organize memory for temporal location information. A Context Boundary Effect (DuBrow & Davachi, 2013) was observed which showed that items within the same temporal context were placed closer together than items which were in adjacent temporal contexts despite being equidistant. This effect did not diminish across trials, suggesting it reflects an underlying organizational principle for the temporal information. Moreover, this contextual information biased the misassignment information such that events were far more likely to be misassigned to the temporal location of the other within-context event than would be expected by chance. Together, this evidence shows the critical interplay of contextual and relational information in the temporal domain and hints at differences in the ability to form relations in the temporal domain when compared to the spatial domain. Additionally, it shows that the influence of contextual information is likely not due to an error per se, but it reflects an organizational principle guiding memory.

In the fourth chapter, the Spatiotemporal Navigation Task that was reexamined through the lens of navigation, rather than reconstruction, in order to determine the degree to which changes in navigation behavior relate to changes in test time performance. Leveraging previous work on spatial navigation (Daugherty et al., 2015; Eichenbaum, 2017c; Kolarik et al., 2017), this chapter provided several measures and models of spatiotemporal navigation of theoretical interest, and introduced a new metric, Lacunarity, reflecting the relative systematicity of navigation. A large number of findings were presented in this chapter, bringing together elements of spatial, temporal, relational, and contextual memory. Improvements in spatial and temporal navigation were shown to relate to improvements in memory in those domains separably, suggesting that spatial and temporal representations may in some way be separable in this task. Relational memory improvements were shown to be tied to changes in navigation complexity and systematicity, suggesting a critical interplay between in-the-moment, memory-guided decision making and subsequent relational memory performance. Evidence was presented suggesting that context boundaries may act as more of a discriminatory feature (at least in this task) than one used to strengthen within-context relational memory organization accuracy; Lastly, a preference towards exploring an otherwise temporally-flexible environment in the implied, forward order with increasing contiguity was shown to be a potentially critical behavior in rapidly improving temporal, relational, and contextual memory organization. Taken together, this chapter established a foundation for using changes in navigation performance to relate to changes in spatial, temporal, relational, and contextual memory organization.

5.2 Organization of Memory Representations

Since the advent of multiple memory systems theory, we've understood that memory cannot be captured uniformly by a single representational abstraction, governed by a single brain region. Through the study of patients with hippocampal damage we have developed an understand that episodic memory (Cohen & Eichenbaum, 1993), navigation (Morris et al., 1982), creativity (Duff et al., 2013), and several other functions (Konkel et al., 2008; Tavares et al., 2015; Warren et al., 2015) are contributed to by hippocampally dependent memory. Relational memory theory (Cohen & Eichenbaum, 1993) has long held that the informational structure of hippocampally dependent memory is underpinned by the binding of relations between arbitrary elements so that these representations can be used flexibly as the environment, task goals, and perceptual cues change. Moreover, these relational representations are thought to be critical for encoding episodic memories (Tulving, 2002) which are characterized by their temporal distinctiveness from other information (i.e. you may park in the same parking garage every day, but today you parked in a particular spot that you must remember in isolation of other episodes related to the spatial location). Some descriptions of hippocampal function go further in specifying that certain subregions underpin particular pattern separation, pattern completion, and binding operations, while cortical interactions allow for the learning of heuristic and statistical information (see Kumaran et al., 2016 for an extensive review). Models of contextual drift (Healey & Kahana, 2016; Howard & Kahana, 2002; Polyn et al., 2009) also assert the criticality of interactions between hippocampal and cortical representations, potentially contributing to the automatic segmentation of events in the service of organizing temporal sequences described by event segmentation theory (Zacks et al., 2007). Additionally, the hippocampus is thought to be part of a complex network, linked with cortical regions, to enable navigation behaviors (Ekstrom et al., 2017) via precision binding (Kolarik et al., 2016, 2017; Yonelinas, 2013) as task demands require (Ekstrom & Ranganath, 2017).

This work sits at the intersection of many of these models and theories by allowing navigation of spatial and temporal information (Chapter 4) and analyzing contextual and relational memory performance (Chapter 3) based on computational methods validated on hippocampally damaged patients (Chapter 2). In agreement with the core ideas of relational memory theory, we show that specific *types* of relations (arbitrary relations that could have just as easily been anything) are impaired with hippocampal damage while spatial locations are not (consistent with the idea that other brain regions may preferential encode spatial location independently of the hippocampus and the fact that place cells exist outside of the hippocampus; see Ekstrom et al., 2017 for a review). However, even once accounting for several types of errors, some residual precision deficits were still present in hippocampal patients, consistent

with the Precision and Binding Model (Yonelinas, 2013). The current work also suggests that, in healthy

adults, differences between spatial and temporal information representations are present, consistent

with suggestions by others that although the hippocampus is capable of binding information within and

across many domains, it may preferentially separate information in spatial and temporal domains

(Eichenbaum, 2017b). Future work using the computational frameworks established in this work, the

Spatiotemporal Navigation Task, and related tasks allowing navigation of alternate domains (visual

perceptual spaces, social spaces, auditory spaces) should examine whether this pattern of results is

present in domains other than space in hippocampal damaged patients.

5.3 Separable vs. Separate Hippocampal Representations

There has been a great deal of debate in the literature over the separateness of different informational

domains in the hippocampus (see Eichenbaum, 2017 for a review). A common perspective is that major

domains such as space and time provide a scaffolding in which binding of information from and to other

domains might be embedded (Ekstrom & Ranganath, 2017). Relatedly, tasks investigating hippocampal

involvement in relations outside of the spatial or temporal domain (in, for instance, social or perceptual

spaces) have recently shown that hippocampal representations may move beyond the constraints of this

scaffolding if task demands require (Tavares et al., 2015; Warren et al., 2015). However, one crucial

question has yet to be completely addressed: are different domains represented entirely separately in

the hippocampus or in some sort of alternate, overlapping set of representations, or something else in-

between? The data here cannot completely disentangle these possibilities, but the techniques

presented here, when combined with neuroimaging and patient behavior studies, may eventually

elucidate the answer to this question. In the Spatiotemporal Navigation Task, differences in relational

memory performance (identity-location misassignments) were seen in space and time, suggesting these

domains may, in fact, be separate. They may be, for instance, multiplexed, with cells switching between

representing spatial and temporal information (see Eichenbaum, 2017 for a review of evidence for this

idea). Alternatively, the representations may overlap insofar as the task demands warrant it. However,

when temporal or spatial information are biased in some way (via attention, task manipulation, reward,

or some other method), the degree of overlap between the domain representations may change. This

hypothesis, that representational separability is dependent on task demands, leads to a wider

suggestion about domain representations in the hippocampus, namely, that *hippocampal relational*

representations are separable but not necessarily separate across domains. This idea is consistent with

relational memory theory (Cohen & Eichenbaum, 1993) but extends the prediction that all arbitrary

relationships can be bound in the hippocampus, to say that the cell firing in association with a domain of

information will depend on demands related to that domain (for example, the goal-directed behavior being targeted by a task). The ability of hippocampal representations to adapt to task demands to optimize the degree of separability of encoding of different variables may be a critical component in memory representations, and as such, the distinction between hippocampal representations across domains being separable, rather than separate, may be of fundamental importance to understanding the structure of hippocampally dependent memory in general.

5.4 Behavioral Inference, Sampling, and Task Richness

Two key contributions of this work are worth reiterating. First, the methodological framework presented in Chapter 2 and utilized in Chapter 3 approaches the issue of understanding information present in a spatial reconstruction from a principled perspective. That is to say, an attempt was made to include as few assumptions as possible about the nature of the high-level structure of the performance and subsequent data in the task, treating it as an information-containing system with no other significant preconceptions. The assumptions are enumerated in that chapter, any of which imply a particular set of choices for analysis. This means that in cases like the comparison of hippocampal damaged patients to comparison participants, adjustments to the assumptions (i.e. the assumption that participants will generally be accurate at placing items back in roughly the right positions) have associated adjustments in the methodology (i.e. using the accuracy threshold from the comparisons, in which the assumption should be valid, as the threshold for the patients). Similarly, other aspects of the method assume certain transformations might be present within the item-item and item-environment relational information (i.e. translation, scaling, and rotation), but the method allows for different transformations to be tested together or in isolation of one-another. Most importantly, by keeping the methodology general, it can be easily applied to other domains (i.e. time), domains with even more dimensions (i.e. 3D space), or even domains with stranger properties (see the Introduction for more information on possible properties of domains). In this way, the tools used to study representations in reconstruction do not have to change even in future explorations that may employ a different set of domains.

The second key contribution is the critical combination of navigation behavior with reconstruction in multiple domains, simultaneously. This allows for three important secondary impacts on the study of memory: 1) The ability to navigate domains other than space makes it possible to examine the impact of exploration and navigation in different domains on subsequent memory. 2) The ability to simultaneously navigate multiple domains means that exploration/navigation behaviors and preferences between domains can be examined (i.e. is there a "hierarchy" of domains and how flexible is such a domain

preference; McKenzie et al., 2014). 3) The overall richness of this task, as well as the resulting findings, shows that many of the effects observed in more restrictive behavioral tasks will still be present when the task is complex, rich, and challenging. This added richness also affords more information and control with which to make observations and predictions about the interplay of different behaviors (navigation and reconstruction) and the brain systems which may support different types of information (scene sampling, contextual information, relational information, and more). In summary, by embracing the rich, complex possibilities available with modern behavioral platforms such as Virtual Reality, we obtain an equal richness and complexity of data in service of enhancing our understanding of the organization of memory.

5.5 Artificial Intelligence and Spatial Reconstructions

One field which is intimately related to the work presented in this document which has yet to be addressed is Machine Learning and Artificial Intelligence. Relational memory theory and contextual organization of memory may become more critical to the development of artificially intelligent (AI) systems in the coming decades as advances in deep learning systems which go beyond reinforcement learning begin to emerge. Spatial reconstruction, as a task, is an incredibly difficult problem for AI systems to perform. In particular, the incredible degree of variability in environments makes navigation in general a challenge. To add on to this the requirement that specific, arbitrary information about elements of the environment need to be remembered and used flexibly at some unknown point in the future makes the task beyond the scope of what current AI can successfully do. Recent advances in deep learning have developed systems which can navigate environments in an end-to-end fashion (i.e. mapping visual input to motor output without any additional "cheating" information provided by the programmer; Mirowski et al., 2017). These systems, however, likely lack the flexible, relational memory representations necessary to perform reconstruction. Some recent work has shown networks which can be trained to understand spatial relational information (Santoro et al., 2017), however, it is unclear how related these representations are to real brain representations (despite the attempts of some authors often make to architect their systems similarly to biological systems at a coarse level). Spatial reconstruction tasks which also involve navigation may serve as a crucial test platform for future AI to determine the degree of flexibility of their representations. Moreover, spatial reconstruction in humans typically involve only one viewing of the information in question. Although so-called "one-shot" learning (Shaban et al., 2017; van der Spoel et al., 2015) is a topic of intense interest in AI, the ability for navigation systems to perform one-shot learning is still lacking. Representational systems which can perform computation similar to what is thought to occur in the hippocampus may serve as one of the

major improvements on deep learning AI in the coming decades, bridging the performance gap on general tasks (and potentially tasks which require creativity) between humans and AI.

5.6 Conclusion

Our understanding of hippocampally-dependent memory organization has come a long way in the past several decades, and with the advent of new technologies and theoretical frameworks, the pace of discovery does not seem to be abating. Here, using a new methodological framework, it has been shown that the hippocampus is necessary for some, but not all, relational information (critically, the arbitrary information but not the location information). These same methods were then applied to a novel task, the Spatiotemporal Navigation Task, which allows for spatial and temporal relational memory and the influence of context to be examined simultaneously. In this task, many more temporal relational memory errors than spatial relational memory errors were present despite the ability to freely navigate time and space, and these errors were heavily influenced by contextual information. The Context Boundary Effect, in which events in time are placed closer together within a context and further apart when separated by a boundary despite being equidistant, was observed in time despite the ability of participants to cross the boundaries in either direction and as often as desired and despite repeated restudy of the same events. Moreover, the relational memory errors were far more common within a context than would be expected by chance, suggesting that the contextual information does not just bias distance judgements but also orders information in a disadvantageous way. Finally, spatiotemporal navigation behavior was examined to determine if particular aspects of navigation in space and time relate to later performance and relational or contextual memory errors. Improvements in spatial and temporal navigation were shown to relate to improvements in memory in those domains separably, suggesting that spatial and temporal representations may in some way be separable in this task. Further, relational memory improvements were shown to be tied to changes in navigation complexity and systematicity, suggesting a critical interplay between in-the-moment, memory-guided decision making and subsequent relational memory efficacy. Additionally, context boundaries were shown to act as more of a discriminatory feature (at least in this task) than one used to strengthen within-context relational memory organization accuracy. Finally, data revealed a preference towards exploring an otherwise temporally-flexible environment in the implied, forward order with increasing contiguity and that exploration pattern was shown to be a potentially critical behavior in rapidly improving temporal, relational, and contextual memory organization.

Together, three main conclusions can be drawn from these data. First, the hippocampus is critical for certain *types* of relational memory information (specifically, the "arbitrary" identity-location information which would be difficult to derive via a heuristic approach), but even with hippocampal damage, memory for the spatial locations of items is not impaired, suggesting this information is being supported by a brain region other than the hippocampus, and critically, not all *types* of relational information require the hippocampus (i.e. some amount of item-item and item-environment information is present). Secondly, temporal and spatial relational memory show key differences in performance (with temporal relational memory errors being far more common than spatial relational memory errors) even when both the spatial location of events and the temporal sequence of events are allowed to be navigated freely, suggesting some degree of separability, perhaps strongly influenced by contextual information, may be present in the representations of different domains of information in hippocampal representations. Moreover, a proof of concept is provided for behavioral paradigms which can assess the behavior surrounding the concept of navigation in domains other than space. And thirdly, changes in navigation behavior in both time and space relate to changes in measures of temporal relational and contextual memory, suggesting that hippocampally dependent relational memory may be critically related to in-the-moment decision making in navigation when the task involves learning of arbitrary event information in multiple domains.

The richness of Spatiotemporal Navigation and Reconstruction has provided (and will likely to continue to provide) interesting findings on the interplay between navigation in space, navigation in time, and how these ultimately may relate to navigation in memory. Through embracing principled approaches to analysis of behavioral data, and the inclusion of complex behavioral mechanics (such as simulated time travel), this work extends our understanding of the role of hippocampal relational memory in overall memory organization.

As a final aside, if any actual time travelers read this document, please feel free to attend my thesis defense at 40.115600°N, -88.227700°W, Friday, April 6th, 2018 CE, 2PM CST (8PM UTC). There will be snacks.

References

Abbott, E. A. (1884). *Flatland: A Romance in Multiple Dimensions*. New York: Dover Thrift.

Aggleton, J. P., & Brown, M. W. (1999). Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. *The Behavioral and Brain Sciences*, *22*(3), 425-44-89. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/11301518

Al-kadi, O. S., & Watson, D. (2008). Texture Analysis of Aggressive and Nonaggressive Lung Tumor CE CT Images, *55*(7), 1822–1830. http://doi.org/10.1109/TBME.2008.919735

Allen, J. S., Tranel, D., Bruss, J., & Damasio, H. (2006). Correlations between regional brain volumes and memory performance in anoxia. *Journal of Clinical and Experimental Neuropsychology*, *28*(4), 457–476. http://doi.org/10.1080/13803390590949287

Allen, R. J., Vargha-Khadem, F., & Baddeley, A. D. (2014). Item-location binding in working memory: Is it hippocampus-dependent? *Neuropsychologia*, *59*(1), 74–84. http://doi.org/10.1016/j.neuropsychologia.2014.04.013

Astur, R. S., Taylor, L. B., Mamelak, A. N., Philpott, L., & Sutherland, R. J. (2002). Humans with hippocampus damage display severe spatial memory impairments in a virtual Morris water task. *Behavioural Brain Research*, *132*(1), 77–84. http://doi.org/10.1016/S0166-4328(01)00399-0

Baker, R., Dexter, M., Hardwicke, T. E., Goldstone, A., & Kourtzi, Z. (2014). Learning to predict: Exposure to temporal sequences facilitates prediction of future events. *Vision Research*, *99*, 124–133. http://doi.org/10.1016/j.visres.2013.10.017

Banta Lavenex, P. A., Colombo, F., Ribordy Lambert, F., & Lavenex, P. (2014). The human hippocampus beyond the cognitive map: evidence from a densely amnesic patient. *Frontiers in Human Neuroscience*, *8*. http://doi.org/10.3389/fnhum.2014.00711

Bartsch, T., Schonfeld, R., Muller, F. J., Alfke, K., Leplow, B., Aldenhoff, J., … Koch, J. M. (2010). Focal Lesions of Human Hippocampal CA1 Neurons in Transient Global Amnesia Impair Place Memory. *Science*, *328*(5984), 1412–1415. http://doi.org/10.1126/science.1188160

Besi, P. J., & Mckay, N. D. (1991). A Method for Registration of 3-D Shapes 1 Introduction 2 Literature Review. In *SPIE Vol. 1611 Sensor Fusion IV* (Vol. 1611, pp. 586–606). http://doi.org/10.1117/12.57955

Birch, A., Osborne, M., & Blunsom, P. (2010). Metrics for MT evaluation: evaluating reordering. *Machine Translation*, *24*(1), 15–26. http://doi.org/10.1007/s10590-009-9066-5

Bonham-Carter, O., Steele, J., & Bastola, D. (2014). Alignment-free genetic sequence comparisons: a review of recent approaches by word analysis. *Briefings in Bioinformatics*, *15*(6), 890–905. http://doi.org/10.1093/bib/bbt052

Brodeur, M. B., Guérard, K., & Bouras, M. (2014). Bank of Standardized Stimuli (BOSS) Phase II: 930 New Normative Photos. *PLoS ONE*, *9*(9), e106953. http://doi.org/10.1371/journal.pone.0106953

Buchanan, T. W. (2005). Emotional Autobiographical Memories in Amnesic Patients with Medial Temporal Lobe Damage. *Journal of Neuroscience*, *25*(12), 3151–3160. http://doi.org/10.1523/JNEUROSCI.4735-04.2005

Bunsey, M., & Eichenbaum, H. (1993). Critical role of the parahippocampal region for paired-associate

learning in rats. *Behavioral Neuroscience*, *107*(5), 740–7. http://doi.org/10.1037/0735-7044.107.5.740

Burkard, R., Dell'Amico, M., & Martello, S. (2012). *Assignment Problems*. Society for Industrial and Applied Mathematics. http://doi.org/10.1137/1.9781611972238

Buzsáki, G. (2005). Theta rhythm of navigation: Link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus*, *15*(7), 827–840. http://doi.org/10.1002/hipo.20113

Carlsson, G. (2009). *Topology and data*. *Bulletin of the American Mathematical Society* (Vol. 46). http://doi.org/10.1090/S0273-0979-09-01249-X

Cederberg, J. N. (2001). Chapter 3.12 Similarity Transformations. In *A Course in Modern Geometries* (pp. 183–189). Springer.

Clark, R., Tahan, A. C., Watson, P. D., Severson, J., Cohen, N. J., & Voss, M. (2017). Aging affects spatial reconstruction more than spatial pattern separation performance even after extended practice. *Hippocampus*, *0*(October 2016), 1–11. http://doi.org/10.1002/hipo.22727

Cline, A. K., & Dhillon, I. S. (2006). Computation of the Singular Value Decomposition. *Handbook of Linear Algebra*, 1–14. http://doi.org/10.1007/BF02251248

Close, C. F. (1902). Map Projections. *The Geographical Journal*, *19*(1), 78. http://doi.org/10.2307/1774978

Cohen, N. J., & Banich, M. T. (2003). Memory. In *Cognitive Neuroscience and Neuropsychology* (2nd ed., pp. 322–364). Boston: Houghton-Mifflin.

Cohen, N. J., & Eichenbaum, H. (1991). The theory that wouldn't die: a critical look at the spatial mapping theory of hippocampal function. *Hippocampus*, *1*(3), 265–268. http://doi.org/10.1002/hipo.450010312

Cohen, N. J., & Eichenbaum, H. (1993). *Memory, Amnesia, and the Hippocampal System.* Cambridge: MIT Press.

Coronel, J. C., Duff, M. C., Warren, D. E., Federmeier, K. D., Gonsalves, B. D., Tranel, D., & Cohen, N. J. (2012). Remembering and voting: Theory and evidence from amnesic patients. *American Journal of Political Science*, *56*(4), 837–848. http://doi.org/10.1111/j.1540-5907.2012.00608.x

Corsi, P. M. (1972). *Human Memory and the Medial Temporal Region of the Brain*. McGill University.

Coxeter, H. S. M. (Harold S. M. (1998). *Non-Euclidean geometry*.

Damasio, A. R. (1989). The Brain Binds Entities and Events by Multiregional Activation from Convergence Zones. *Neural Computation*, *1*(1), 123–132. http://doi.org/10.1162/neco.1989.1.1.123

Daugherty, A. M., Yuan, P., Dahle, C. L., Bender, A. R., Yang, Y., & Raz, N. (2015). Path complexity in virtual water maze navigation: Differential associations with age, sex, and regional brain volume. *Cerebral Cortex*, *25*(9), 3122–3131. http://doi.org/10.1093/cercor/bhu107

Davachi, L., & DuBrow, S. (2015). How the hippocampus preserves order: the role of prediction and context. *Trends in Cognitive Sciences*, *19*(2), 92–99. http://doi.org/10.1016/j.tics.2014.12.004

Davis, C. P., Franklin, L. M., Johnson, G. S., & Schrott, L. M. (2010). Prenatal oxycodone exposure impairs spatial learning and/or memory in rats. *Behavioural Brain Research*, *212*(1), 27–34.

http://doi.org/10.1016/j.bbr.2010.03.022

Dede, A. J. O., Frascino, J. C., Wixted, J. T., & Squire, L. R. (2016). Learning and remembering real-world events after medial temporal lobe damage. *Proceedings of the National Academy of Sciences*, *113*(47), 13480–13485. http://doi.org/10.1073/pnas.1617025113

Dinu, L. P., & Ionescu, R. (2012). An Efficient Rank Based Approach for Closest String and Closest Substring. *PLoS ONE*, *7*(6), e37576. http://doi.org/10.1371/journal.pone.0037576

DuBrow, S., & Davachi, L. (2013). The influence of context boundaries on memory for the sequential order of events. *Journal of Experimental Psychology. General*, *142*(4), 1277–86. http://doi.org/10.1037/a0034024

DuBrow, S., & Davachi, L. (2016). Temporal binding within and across events. *Neurobiology of Learning and Memory*, *134*, 107–114. http://doi.org/10.1016/j.nlm.2016.07.011

Duchowski, A. T., Shivashankaraiah, V., Rawls, T., Gramopadhye, A. K., Melloy, B. J., & Kanki, B. (2000). Binocular eye tracking in virtual reality for inspection training. In *Proceedings of the symposium on Eye tracking research & applications - ETRA '00* (pp. 89–96). New York, New York, USA: ACM Press. http://doi.org/10.1145/355017.355031

Dudchenko, P. (2010). *Why People Get Lost : The Psychology and Neuroscience of Spatial Cognition Abstract and Keywords Taxonomies of wayfinding*. http://doi.org/10.1093/acprof

Duff, M. C., Kurczek, J., Rubin, R., Cohen, N. J., & Tranel, D. (2013). Hippocampal amnesia disrupts creative thinking. *Hippocampus*, *23*(12), 1143–1149. http://doi.org/10.1002/hipo.22208

Dutta, A., & Nairne, J. S. (1993). The separability of space and time: dimensional interaction in the memory trace. *Memory & Cognition*, *21*(4), 440–8. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8350735

Eichenbaum, H. (2004). Hippocampus: Cognitive processes and neural representations that underlie declarative memory. *Neuron*, *44*(1), 109–120. http://doi.org/10.1016/j.neuron.2004.08.028

Eichenbaum, H. (2013). Memory on time. *Trends in Cognitive Sciences*, *17*(2), 88. http://doi.org/10.1016/j.tics.2012.12.007

Eichenbaum, H. (2014). Time cells in the hippocampus: a new dimension for mapping memories. *Nature Reviews. Neuroscience*, *15*(11), 732–44. http://doi.org/10.1038/nrn3827

Eichenbaum, H. (2015). The Hippocampus as a Cognitive Map ... of Social Space. *Neuron*, *87*(1), 9–11. http://doi.org/10.1016/j.neuron.2015.06.013

Eichenbaum, H. (2016). Memory: Organization and Control. *Annual Review of Psychology*, (September), 1–27. http://doi.org/10.1146/annurev-psych-010416-044131

Eichenbaum, H. (2017a). Barlow versus Hebb: When is it time to abandon the notion of feature detectors and adopt the cell assembly as the unit of cognition? *Neuroscience Letters*. http://doi.org/10.1016/j.neulet.2017.04.006

Eichenbaum, H. (2017b). On the Integration of Space, Time, and Memory. *Neuron*, *95*(5), 1007–1018. http://doi.org/10.1016/j.neuron.2017.06.036

Eichenbaum, H. (2017c). The role of the hippocampus in navigation is memory. *Journal of*

Neurophysiology, *117*(4), 1785–1796. http://doi.org/10.1152/jn.00005.2017

Eichenbaum, H. (2017d). Time (and space) in the hippocampus. *Current Opinion in Behavioral Sciences*, *17*, 65–70. http://doi.org/10.1016/j.cobeha.2017.06.010

Eichenbaum, H., & Cohen, N. J. (2001). *From Conditioning to Conscious Recollection: Memory Systems of the Brain*. New York: Oxford University Press.

Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., & Tanila, H. (1999). The Hippocampus, Memory, Review and Place Cells: Is It Spatial Memory or a Memory Space? *Neuron*, *23*, 209–226. http://doi.org/10.1016/S0896-6273(00)80773-4

Eichenbaum, H., & Fortin, N. J. (2009). The neurobiology of memory based predictions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1521), 1183–1191. http://doi.org/10.1098/rstb.2008.0306

Ekstrom, A. D., Huffman, D. J., & Starrett, M. (2017). Interacting networks of brain regions underlie human spatial navigation: a review and novel synthesis of the literature. *Journal of Neurophysiology*, *118*(6), 3328–3344. http://doi.org/10.1152/jn.00531.2017

Ekstrom, A. D., & Isham, E. A. (2017). Human spatial navigation: representations across dimensions and scales. *Current Opinion in Behavioral Sciences*, *17*, 84–89. http://doi.org/10.1016/j.cobeha.2017.06.005

Ekstrom, A. D., & Ranganath, C. (2017). Space, time, and episodic memory: The hippocampus is all over the cognitive map. *Hippocampus*, *93052*, 1–16. http://doi.org/10.1002/hipo.22750

Epstein, R. A. (2014). Neural Systems for Visual Scene Recognition. *Scene Vision*, 105–134. http://doi.org/10.7551/mitpress/9780262027854.001.0001

Ezzyat, Y., & Davachi, L. (2014). Similarity breeds proximity: Pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity. *Neuron*, *81*(5), 1179–1189. http://doi.org/10.1016/j.neuron.2014.01.042

Fortin, N. J., Agster, K. L., & Eichenbaum, H. B. (2002). Critical role of the hippocampus in memory for sequences of events. *Nature Neuroscience*, *5*(5), 458–462. http://doi.org/10.1038/nn834

Franz, M. O., & Mallot, H. A. (2000). Biomimetic robot navigation. *Robotics and Autonomous Systems*, *30*(1–2), 133–153. http://doi.org/10.1016/S0921-8890(99)00069-X

Gallistel, C. R. (1990). *The organization of learning*. *MIT Press series in learning, development, and conceptual change*.

Garcia, F., Carrère, P., Soussana, J. F., & Baumont, R. (2005). Characterisation by fractal analysis of foraging paths of ewes grazing heterogeneous swards. *Applied Animal Behaviour Science*, *93*(1–2), 19–37. http://doi.org/10.1016/j.applanim.2005.01.001

Gautestad, A. O. (2011). Memory matters: Influence from a cognitive map on animal space use. *Journal of Theoretical Biology*, *287*, 26–36. http://doi.org/10.1016/j.jtbi.2011.07.010

Gautestad, A. O., & Mysterud, I. (2010). Spatial memory, habitat auto-facilitation and the emergence of fractal home range patterns. *Ecological Modelling*, *221*(23), 2741–2750. http://doi.org/10.1016/j.ecolmodel.2010.08.014

Gelman, A., & Hill, J. (2007). *Data Analysis Using Regression and Multilevel/Hierarchical Models.* New York: Cambridge University Press.

Geva-Sagiv, M., Las, L., Yovel, Y., & Ulanovsky, N. (2015). Spatial cognition in bats and rats: from sensory acquisition to multiscale maps and navigation. *Nature Reviews Neuroscience*, *16*(2), 94–108. http://doi.org/10.1038/nrn3888

Ghrist, R. (2007). Barcodes: The persistent topology of data. *Bulletin of the American Mathematical Society*, *45*(1), 61–76. http://doi.org/10.1090/S0273-0979-07-01191-3

Giancarlo, R., Scaturro, D., & Utro, F. (2009). Textual data compression in computational biology: a synopsis. *Bioinformatics*, *25*(13), 1575–1586. http://doi.org/10.1093/bioinformatics/btp117

Gilbert, P. E., Kesner, R. P., & DeCoteau, W. E. (1998). Memory for spatial location: role of the hippocampus in mediating spatial pattern separation. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *18*(2), 804–10. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/9425021

Giovanello, K. S., Verfaellie, M., & Keane, M. M. (2003). Disproportionate deficit in associative recognition relative to item recognition in global amnesia. *Cognitive, Affective & Behavioral Neuroscience*, *3*(3), 186–94. http://doi.org/10.3758/CABN.3.3.186

Golub, G., & Kahan, W. (1965). CALCULATING THE SINGULAR VALUES AND PSEUDO-INVERSE OF A MATRIX. *J. SIAM Numer, Anal, Ser. B*, *2*(2), 205–224.

Gothard, K. M., Skaggs, W. E., & McNaughton, B. L. (1996). Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *16*(24), 8027–40. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8987829

Gothard, K. M., Skaggs, W. E., Moore, K. M., & McNaughton, B. L. (1996). Binding of hippocampal CA1 neural activity to multiple reference frames in a landmark-based navigation task. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *16*(2), 823–35. http://doi.org/0270-6474/96/160823-13

Hannula, D. E. (2010). Worth a glance: using eye movements to investigate the cognitive neuroscience of memory. *Frontiers in Human Neuroscience*, *4*(October), 1–16. http://doi.org/10.3389/fnhum.2010.00166

Hannula, D. E., & Ranganath, C. (2009). The Eyes Have It: Hippocampal Activity Predicts Expression of Memory in Eye Movements. *Neuron*, *63*(5), 592–599. http://doi.org/10.1016/j.neuron.2009.08.025

Hannula, D. E., Ryan, J. D., Tranel, D., & Cohen, N. J. (2007). Rapid Onset Relational Memory Effects Are Evident in Eye Movement Behavior, but Not in Hippocampal Amnesia. *Journal of Cognitive Neuroscience*, *19*(10), 1690–1705. http://doi.org/10.1162/jocn.2007.19.10.1690

Hannula, D. E., Tranel, D., & Cohen, N. J. (2006). The Long and the Short of It: Relational Memory Impairments in Amnesia, Even at Short Lags. *Journal of Neuroscience*, *26*(32), 8352–8359. http://doi.org/10.1523/JNEUROSCI.5222-05.2006

Hartley, T., Bird, C. M., Chan, D., Cipolotti, L., Husain, M., Vargha-Khadem, F., & Burgess, N. (2007). The hippocampus is required for short-term topographical memory in humans. *Hippocampus*, *17*(1), 34–48. http://doi.org/10.1002/hipo.20240

Hartley, T., Lever, C., Burgess, N., & O'Keefe, J. (2013). Space in the brain: how the hippocampal formation supports spatial cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1635), 20120510–20120510. http://doi.org/10.1098/rstb.2012.0510

Hartley, T., Trinkler, I., & Burgess, N. (2004). Geometric determinants of human spatial memory. *Cognition*, *94*(1), 39–75. http://doi.org/10.1016/j.cognition.2003.12.001

Hasselmo, M. E., & Howard Eichenbaum. (2005). Hippocampal mechanisms for the context-dependent retrieval of episodes. *Neural Networks*, *18*(9), 1172–1190. http://doi.org/10.1016/j.neunet.2005.08.007

Hawking, S. W., & Ellis, G. F. R. (1973). *The Large Scale Structure of Space–Time*. *Book*. http://doi.org/10.1017/CBO9780511524646

Hayes, S. M., Ryan, L., Schnyer, D. M., & Nadel, L. (2004). An fMRI Study of Episodic Memory: Retrieval of Object, Spatial, and Temporal Information. *Behavioral Neuroscience*, *118*(5), 885–896. http://doi.org/10.1037/0735-7044.118.5.885

Healey, M. K., & Kahana, M. J. (2016). A four-component model of age-related memory change. *Psychological Review*, *123*(1), 23–69. http://doi.org/10.1037/rev0000015

Heusser, A. C., Ezzyat, Y., Shiff, I., & Davachi, L. (2018). Perceptual boundaries cause mnemonic trade-offs between local boundary processing and across-trial associative binding. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, (212), 1–48.

Hopcroft, J., & Tarjan, R. (1973). Algorithm 447: efficient algorithms for graph manipulation. *Communications of the ACM*, *16*(6), 372–378. http://doi.org/10.1145/362248.362272

Horner, A. J., Bisby, J. A., Wang, A., Bogus, K., & Burgess, N. (2016). The role of spatial boundaries in shaping long-term event representations. *Cognition*, *154*, 151–164. http://doi.org/10.1016/j.cognition.2016.05.013

Howard, M. W., Fotedar, M. S., Datey, A. V, & Hasselmo, M. E. (2005). The Temporal Context Model in spatial navigation and relational learning: Toward a common explanation of medial temporal lobe function across domains. *Psychol Rev*, *112*(1), 75–116. http://doi.org/10.1037/0033-295X.112.1.75

Howard, M. W., & Kahana, M. J. (2002). A Distributed Representation of Temporal Context. *Journal of Mathematical Psychology*, *46*(3), 269–299. http://doi.org/10.1006/jmps.2001.1388

Howard, M. W., MacDonald, C. J., Tiganj, Z., Shankar, K. H., Du, Q., Hasselmo, M. E., & Eichenbaum, H. (2014). A Unified Mathematical Framework for Coding Time, Space, and Sequences in the Hippocampal Region. *Journal of Neuroscience*, *34*(13), 4692–4707. http://doi.org/10.1523/JNEUROSCI.5808-12.2014

Howard, M. W., & Natu, V. S. (2005). Place from time: Reconstructing position from a distributed representation of temporal context. *Neural Networks*, *18*(9), 1150–1162. http://doi.org/10.1016/j.neunet.2005.08.002

Howard, M. W., Shankar, K. H., Aue, W. R., & Criss, A. H. (2015). A distributed representation of internal time. *Psychological Review*, *122*(1), 24–53. http://doi.org/10.1037/a0037840

Howard, M. W., Viskontas, I. V, Shankar, K. H., & Fried, I. (2012). Ensembles of human MTL neurons "jump back in time" in response to a repeated stimulus. *Hippocampus*, *22*(9), 1833–47.

http://doi.org/10.1002/hipo.22018

Hsieh, L.-T., Gruber, M. J., Jenkins, L. J., & Ranganath, C. (2014). Hippocampal Activity Patterns Carry Information about Objects in Temporal Context. *Neuron*, *81*(5), 1165–1178. http://doi.org/10.1016/j.neuron.2014.01.015

Hsieh, L. T., & Ranganath, C. (2015). Cortical and subcortical contributions to sequence retrieval: Schematic coding of temporal context in the neocortical recollection network. *NeuroImage*, *121*(121), 78–90. http://doi.org/10.1016/j.neuroimage.2015.07.040

Huth, J. E. (2013). *The Lost Art of Finding Our Way*. Cambridge, MA and London, England: Harvard University Press. http://doi.org/10.4159/harvard.9780674074811

Huttenlocher, J., & Presson, C. C. (1979). The coding and transformation of spatial information. *Cognitive Psychology*, *11*(3), 375–394. http://doi.org/10.1016/0010-0285(79)90017-3

Insausti, R., Amaral, D. G., & Cowan, W. M. (1987). The entorhinal cortex of the monkey: II. Cortical afferents. *The Journal of Comparative Neurology*, *264*(3), 356–395. http://doi.org/10.1002/cne.902640306

Jelinek, H. F., Karperien, A., & Milosevic, N. T. (2011). Reviewing Lacunarity Analysis and Classification of Microglia in Neuroscience. *Proceedings of the 8th European Conference on Mathematical and Theoretical Biology*, *88*.

Jeneson, A., Mauldin, K. N., & Squire, L. R. (2010). Intact Working Memory for Relational Information after Medial Temporal Lobe Damage. *Journal of Neuroscience*, *30*(41), 13624–13629. http://doi.org/10.1523/JNEUROSCI.2895-10.2010

Julian, J. B., Ryan, J., Hamilton, R. H., & Epstein, R. A. (2016). The Occipital Place Area Is Causally Involved in Representing Environmental Boundaries during Navigation. *Current Biology*, *26*(8), 1104–1109. http://doi.org/10.1016/j.cub.2016.02.066

Juliani, A. (2016). Simple Reinforcement Learning with Tensorflow Part 8: Asynchronous Actor-Critic Agents (A3C). Retrieved from https://medium.com/emergent-future/simple-reinforcement-learning-with-tensorflow-part-8-asynchronous-actor-critic-agents-a3c-c88f72a5e9f2

Kahana, M. J., & Howard, M. W. (2005). Spacing and lag effects in free recall of pure lists. *Psychonomic Bulletin & Review*, *12*(1), 159–164. http://doi.org/10.3758/BF03196362

Kesner, R. P., Gilbert, P. E., & Barua, L. A. (2002). The role of the hippocampus in memory for the temporal order of a sequence of odors. *Behavioral Neuroscience*, *116*(2), 286–290. http://doi.org/10.1037/0735-7044.116.2.286

Kessels, R. P. C., van Zandvoort, M. J. E., Postma, A., Kappelle, L. J., & de Haan, E. H. F. (2000). The Corsi Block-Tapping Task: Standardization and Normative Data. *Applied Neuropsychology*, *7*(4), 252–258. http://doi.org/10.1207/S15324826AN0704_8

Knierim, J. J., & Neunuebel, J. P. (2016). Tracking the flow of hippocampal computation: Pattern separation, pattern completion, and attractor dynamics. *Neurobiology of Learning and Memory*, *129*, 38–49. http://doi.org/10.1016/j.nlm.2015.10.008

Kolarik, B. S., Baer, T., Shahlaie, K., Yonelinas, A. P., & Ekstrom, A. D. (2017). Close but no cigar: Spatial precision deficits following medial temporal lobe lesions provide novel insight into theoretical

models of navigation and memory. *Hippocampus*, (August), 1–11. http://doi.org/10.1002/hipo.22801

Kolarik, B. S., Shahlaie, K., Hassan, A., Borders, A. A., Kaufman, K. C., Gurkoff, G., … Ekstrom, A. D. (2016). Impairments in precision, rather than spatial strategy, characterize performance on the virtual Morris Water Maze: A case study. *Neuropsychologia*, *80*, 90–101. http://doi.org/10.1016/j.neuropsychologia.2015.11.013

Konkel, A. (2009). Relational memory and the hippocampus: Representations and methods. *Frontiers in Neuroscience*, *3*(2), 166–174. http://doi.org/10.3389/neuro.01.023.2009

Konkel, A., Warren, D. E., Duff, M. C., Tranel, D. N., & Cohen, N. J. (2008). Hippocampal amnesia impairs all manner of relational memory. *Frontiers in Human Neuroscience*, *2*(October), 15. http://doi.org/10.3389/neuro.09.015.2008

Kraus, B. J. J., Brandon, M. P. P., Robinson, R. J. J., Connerney, M. A. A., Hasselmo, M. E. E., & Eichenbaum, H. (2015). During Running in Place, Grid Cells Integrate Elapsed Time and Distance Run. *Neuron*, *88*(3), 578–589. http://doi.org/10.1016/j.neuron.2015.09.031

Kraus, B., Robinson, R., White, J., Eichenbaum, H., & Hasselmo, M. (2013). Hippocampal "Time Cells": Time versus Path Integration. *Neuron*, *78*(6), 1090–1101. http://doi.org/10.1016/j.neuron.2013.04.015

Kropff, E., Carmichael, J. E., Moser, M.-B., & Moser, E. I. (2015). Speed cells in the medial entorhinal cortex. *Nature*, *523*(7561), 419–424. http://doi.org/10.1038/nature14622

Kuhn, H. W. (1956). Variants of the hungarian method for assignment problems. *Naval Research Logistics Quarterly*, *3*(4), 253–258. http://doi.org/10.1002/nav.3800030404

Kuhn, H. W. (2010). The Hungarian Method for the Assignment Problem. In *50 Years of Integer Programming 1958-2008* (pp. 29–47). Berlin, Heidelberg, Heidelberg: Springer Berlin Heidelberg. http://doi.org/10.1007/978-3-540-68279-0_2

Kuliga, S. F., Thrash, T., Dalton, R. C., & Hölscher, C. (2015). Virtual reality as an empirical research tool - Exploring user experience in a real building and a corresponding virtual model. *Computers, Environment and Urban Systems*, *54*, 363–375. http://doi.org/10.1016/j.compenvurbsys.2015.09.006

Kumaran, D., Hassabis, D., & McClelland, J. L. (2016). What Learning Systems do Intelligent Agents Need? Complementary Learning Systems Theory Updated. *Trends in Cognitive Sciences*, *20*(7), 512–534. http://doi.org/10.1016/j.tics.2016.05.004

Kveraga, K., & Bar, M. (2015). *Scene Vision: Making Sense of What We See*. http://doi.org/10.1177/0301006615600645

Lee, A. C. H., Buckley, M. J., Pegman, S. J., Spiers, H., Scahill, V. L., Gaffan, D., … Graham, K. S. (2005). Specialization in the medial temporal lobe for processing of objects and scenes. *Hippocampus*, *15*(6), 782–797. http://doi.org/10.1002/hipo.20101

Lenck-Santini, P.-P., Save, E., & Poucet, B. (2001). Place-cell firing does not depend on the direction of turn in a Y-maze alternation task. *European Journal of Neuroscience*, *13*(5), 1055–1058. http://doi.org/10.1046/j.0953-816x.2001.01481.x

Leutgeb, J. K., Leutgeb, S., Moser, M.-B., & Moser, E. I. (2007). Pattern Separation in the Dentate Gyrus and CA3 of the Hippocampus. *Science*, *315*(5814), 961–966. http://doi.org/10.1126/science.1135801

Levitt, T. S., & Lawton, D. T. (1990). Qualitative navigation for mobile robots. *Artificial Intelligence*, *44*(3), 305–360. http://doi.org/10.1016/0004-3702(90)90027-W

Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*, *6*(6), 579–90. http://doi.org/10.1002/(SICI)1098-1063(1996)6:6<579::AID-HIPO3>3.0.CO;2-C

Libby, L. A., Hannula, D. E., & Ranganath, C. (2014). Medial Temporal Lobe Coding of Item and Spatial Information during Relational Binding in Working Memory. *Journal of Neuroscience*, *34*(43), 14233–14242. http://doi.org/10.1523/JNEUROSCI.0655-14.2014

Lisman, J., & Redish, A. D. (2009). Prediction, sequences and the hippocampus. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1521), 1193–1201. http://doi.org/10.1098/rstb.2008.0316

Lucas, H. D., Duff, M. C., & Cohen, N. J. (2017). The hippocampus promotes effective saccadic information gathering in humans. In *Cognitive Neuroscience Society 24th Annual Meeting*. San Francisco.

Lucas, H. D., Monti, J. M., Mcauley, E., Watson, P. D., Kramer, A. F., & Cohen, N. J. (2016). Relational Memory and Self-Efficacy Measures Reveal Distinct Profiles of Subjective Memory Concerns in Older Adults. *Neuropsychology* , *30*(5), 568–578. http://doi.org/10.1037/neu0000275

Luft, C. D. B., Meeson, A., Welchman, A. E., & Kourtzi, Z. (2015). Decoding the future from past experience: learning shapes predictions in early visual cortex. *Journal of Neurophysiology*, *113*(9), 3159–3171. http://doi.org/10.1152/jn.00753.2014

M.Cover, T. (2006). *Elements Of Information Theory* (2nd ed.). Wiley. Retrieved from http://books.google.fr/books?id=VWq5GG6ycxMC&printsec=frontcover&dq=intitle:Elements+of+Information+Theory&hl=&cd=1&source=gbs_api%5Cnpapers2://publication/uuid/BAF426F8-5A4F-44A4-8333-FA8187160D9B

Maguire, E. A., Burgess, N., O'Keefe, J., & O'Keefe, J. (1999). Human spatial navigation: cognitive maps, sexual dimorphism, and neural substrates. *Current Opinion in Neurobiology*, *9*(2), 171–7. http://doi.org/10.1016/S1474-4422(08)70216-3

Mandelbrot, B. (1967). How Long Is the Coast of Britain? Statistical Self-Similarity and Fractional Dimension. *Science*, *156*(3775), 636–638. http://doi.org/10.1126/science.156.3775.636

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. *Psychological Review*, *102*(3), 419–457. http://doi.org/10.1037/0033-295X.102.3.419

McHugh, T. J., Jones, M. W., Quinn, J. J., Balthasar, N., Coppari, R., Elmquist, J. K., … Tonegawa, S. (2007). Dentate Gyrus NMDA Receptors Mediate Rapid Pattern Separation in the Hippocampal Network. *Science*, *317*(5834), 94–99. http://doi.org/10.1126/science.1140263

McKenzie, S., Frank, A. J. J., Kinsky, N. R. R., Porter, B., Rivière, P. D. D., & Eichenbaum, H. (2014).

Hippocampal Representation of Related and Opposing Memories Develop within Distinct, Hierarchically Organized Neural Schemas. *Neuron*, *83*(1), 202–215. http://doi.org/10.1016/j.neuron.2014.05.019

McNaughton, B. L., Barnes, C. A., & O'Keefe, J. (1983). The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Experimental Brain Research*, *52*(1), 41–9. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/6628596

Mirowski, P., Pascanu, R., Viola, F., Soyer, H., Ballard, A. J., Banino, A., … Hadsell, R. (2016). Learning to Navigate in Complex Environments. Retrieved from http://arxiv.org/abs/1611.03673

Moffat, S. D. (2009). Aging and spatial navigation: What do we know and where do we go? *Neuropsychology Review*, *19*(4), 478–489. http://doi.org/10.1007/s11065-009-9120-3

Moffat, S. D., & Resnick, S. M. (2002). Effects of age on virtual environment place navigation and allocentric cognitive mapping. *Behavioral Neuroscience*, *116*(5), 851–9. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12369805

Monti, J. M., Cooke, G. E., Watson, P. D., Voss, M. W., Kramer, A. F., & Cohen, N. J. (2015). Relating Hippocampus to Relational Memory Processing across Domains and Delays. *Journal of Cognitive Neuroscience*, *27*(2), 234–245. http://doi.org/10.1162/jocn_a_00717

Morris, R. G. M. (1981). Spatial localization does not require the presence of local cues. *Learning and Motivation*, *12*(2), 239–260. http://doi.org/10.1016/0023-9690(81)90020-5

Morris, R. G. M., Garrud, P., Rawlins, J. N. P., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*, *297*(5868), 681–683. http://doi.org/10.1038/297681a0

Muller, R. U., & Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *7*(7), 1951–68. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/3612226

Munkres, J. (1957). Algorithms for the Assignment and Transportation Problems. *Journal of the Society for Industrial and Applied Mathematics*, *5*(1), 32–38. http://doi.org/10.1137/0105003

Murdock, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology*, *64*(5), 482–488. http://doi.org/10.1037/h0045106

Murray, E. A., Gaffan, D., & Mishkin, M. (1993). Neural substrates of visual stimulus-stimulus association in rhesus monkeys. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *13*(10), 4549–61. http://doi.org/0270-647/93/134549-13$05.00/0

Nakazawa, K., Quirk, M. C., Chitwood, R. A. R. A., Watanabe, M., Yeckel, M. F. M. F., Sun, L. D. L. D., … Tonegawa, S. L. B. science. 297. 2002. 21. (2002). Requirement for Hippocampal CA3 NMDA Receptors in Associative Memory Recall. *Science*, *297*(5579), 211–218. http://doi.org/10.1126/science.1071795.Requirement

Naya, Y., & Suzuki, W. A. (2011). Integrating What and When Across the Primate Medial Temporal Lobe. *Science*, *333*(6043), 773–776. http://doi.org/10.1126/science.1206773

Neunuebel, J. P., & Knierim, J. J. (2014). CA3 Retrieves Coherent Representations from Degraded Input: Direct Evidence for CA3 Pattern Completion and Dentate Gyrus Pattern Separation. *Neuron*, *81*(2), 416–427. http://doi.org/10.1016/j.neuron.2013.11.017

Nielson, D. M., Smith, T. A., Sreekumar, V., Dennis, S., & Sederberg, P. B. (2015). Human hippocampus represents space and time during retrieval of real-world memories. *Proceedings of the National Academy of Sciences*, *112*(35), 11078–11083. http://doi.org/10.1073/pnas.1507104112

Noether, E. (1971). Invariant variation problems. *Transport Theory and Statistical Physics*, *1*(3), 186–207. http://doi.org/10.1080/00411457108231446

O'Keefe, J. (1991). An allocentric spatial model for the hippocampal cognitive map. *Hippocampus*, *1*(3), 230–235. http://doi.org/10.1002/hipo.450010303

O'Keefe, J., Burgess, N., O' Keefe, J., Burgess, N., O'Keefe, J., & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, *381*(6581), 425–428. http://doi.org/10.1038/381425a0

O'Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, *34*(1), 171–175. http://doi.org/10.1016/0006-8993(71)90358-1

O'Keefe, J., & Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford: Oxford University Press.

O'Reilly, R. C., & Rudy, J. W. (2001). Conjunctive Representations in Learning and Memory: Principles of Cortical and Hippocampal Function. *Psychological Review*, *108*(1), 83–95. http://doi.org/10.1037//0033-295X.

Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsaki, G. (2008). Internally Generated Cell Assembly Sequences in the Rat Hippocampus. *Science*, *321*(5894), 1322–1327. http://doi.org/10.1126/science.1159775

Paz, R., Gelbard-Sagiv, H., Mukamel, R., Harel, M., Malach, R., & Fried, I. (2010). A neural substrate in the human hippocampus for linking successive events. *Proceedings of the National Academy of Sciences*, *107*(13), 6046–6051. http://doi.org/10.1073/pnas.0910834107

Pecchia, T., & Vallortigara, G. (2010). View-based strategy for reorientation by geometry. *Journal of Experimental Biology*, *213*(17), 2987–2996. http://doi.org/10.1242/jeb.043315

Perrett, D. I., & Oram, M. W. (1993). Neurophysiology of shape processing. *Image and Vision Computing*, *11*(6), 317–333. http://doi.org/10.1016/0262-8856(93)90011-5

Plotnick, R. E., Gardner, R. H., Hargrove, W. W., Prestegaard, K., & Perlmutter, M. (1996). Lacunarity analysis: A general technique for the analysis of spatial patterns. *Physical Review E*, *53*(5), 5461–5468. http://doi.org/10.1103/PhysRevE.53.5461

Polyn, S. M., Norman, K. a, & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, *116*(1), 129–156. http://doi.org/10.1037/a0014420

Poucet, B. (1993). Spatial cognitive maps in animals: New hypotheses on their structure and neural mechanisms. *Psychological Review*, *100*(2), 163–182. http://doi.org/10.1037/0033-295X.100.2.163

Ranck Jr, J. B. (1984). Head direction cells in the deep cell layer of dorsal presubiculum in freely moving rats. *Soc Neurosci Abstr.*, *10*(176), 12.

Ranganath, C., & Blumenfeld, R. S. (2005). Doubts about double dissociations between short- and long-term memory. *Trends in Cognitive Sciences*, *9*(8), 374–380.

http://doi.org/10.1016/j.tics.2005.06.009

Ranganath, C., & Blumenfeld, R. S. (2008). Prefrontal Cortex and Memory. In J. H. Byrne (Ed.), *Learning and Memory: A Comprehensive Reference* (3rd ed., pp. 261–279). Oxford: Acadmic Press.

Ranganath, C., & D'Esposito, M. (2001). Medial Temporal Lobe Activity Associated with Active Maintenance of Novel Information. *Neuron*, *31*(5), 865–873. http://doi.org/10.1016/S0896-6273(01)00411-1

Ranganath, C., & Hsieh, L. T. (2016). The hippocampus: A special place for time. *Annals of the New York Academy of Sciences*, *1369*(1), 93–110. http://doi.org/10.1111/nyas.13043

Ranganath, C., & Ritchey, M. (2012a). Two cortical systems for memory- guided behaviour. *Nature Publishing Group*, *13*(10), 713–726. http://doi.org/10.1038/nrn3338

Ranganath, C., & Ritchey, M. (2012b). Two cortical systems for memory-guided behaviour. *Nature Reviews Neuroscience*, *13*(10), 713–726. http://doi.org/10.1038/nrn3338

Rolls, E. T., & Kesner, R. P. (2006). A computational theory of hippocampal function, and empirical tests of the theory. *Progress in Neurobiology*, *79*(1), 1–48. http://doi.org/10.1016/j.pneurobio.2006.04.005

Rondina, R., Curtiss, K., Meltzer, J. A., Barense, M. D., & Ryan, J. D. (2016). The organisation of spatial and temporal relations in memory. *Memory*, *8211*(March), 1–14. http://doi.org/10.1080/09658211.2016.1182553

Ryan, J. D., & Cohen, N. J. (2004). Processing and short-term retention of relational information in amnesia. *Neuropsychologia*, *42*(4), 497–511. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/14728922

Ryan, L., Lin, C.-Y., Ketcham, K., & Nadel, L. (2009). The role of medial temporal lobe in retrieving spatial and nonspatial relations from episodic and semantic memory. *Hippocampus*, NA-NA. http://doi.org/10.1002/hipo.20607

Santoro, A., Raposo, D., Barrett, D. G. T., Malinowski, M., Pascanu, R., Battaglia, P., & Lillicrap, T. (2017). A simple neural network module for relational reasoning, 1–16. Retrieved from http://arxiv.org/abs/1706.01427

Sarel, A., Finkelstein, A., Las, L., & Ulanovsky, N. (2017). Vectorial representation of spatial goals in the hippocampus of bats. *Science*, *355*(6321), 176–180. http://doi.org/10.1126/science.aak9589

Sargent, J. Q., Zacks, J. M., Hambrick, D. Z., Zacks, R. T., Kurby, C. A., Bailey, H. R., … Beck, T. M. (2013). Event segmentation ability uniquely predicts event memory. *Cognition*, *129*(2), 241–255. http://doi.org/10.1016/j.cognition.2013.07.002

Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: a neural network modelling approach to reconciling episodic memory with statistical learning. *bioRxiv*, *372*(1711), 51870. http://doi.org/10.1101/051870

Schiller, D., Eichenbaum, H., Buffalo, E. A., Davachi, L., Foster, D. J., Leutgeb, S., & Ranganath, C. (2015). Memory and Space: Towards an Understanding of the Cognitive Map. *Journal of Neuroscience*, *35*(41), 13904–13911. http://doi.org/10.1523/JNEUROSCI.2618-15.2015

Schultheis, M. T., Himelstein, J., & Rizzo, A. a. (2002). Virtual reality and neuropsychology: upgrading the

current tools. *The Journal of Head Trauma Rehabilitation*, *17*(5), 378–394. http://doi.org/10.1097/00001199-200210000-00002

Schwarb, H., Johnson, C. L., Daugherty, A. M., Hillman, C. H., Kramer, A. F., Cohen, N. J., & Barbey, A. K. (2017). Aerobic fitness, hippocampal viscoelasticity, and relational memory performance. *NeuroImage*, *153*(March), 179–188. http://doi.org/10.1016/j.neuroimage.2017.03.061

Schwarb, H., Johnson, C. L., McGarry, M. D. J., & Cohen, N. J. (2016). Medial temporal lobe viscoelasticity and relational memory performance. *NeuroImage*, *132*, 534–541. http://doi.org/10.1016/j.neuroimage.2016.02.059

Sederberg, P. B., Howard, M. W., & Kahana, M. J. (2008). A context-based theory of recency and contiguity in free recall. *Psychological Review*, *115*(4), 893–912. http://doi.org/10.1037/a0013396

Sei, Y., Arora, P. K., Skolnick, P., & Paul, I. A. (1992). Spatial learning impairment in a murine model of AIDS. *FASEB Journal*, *6*(11).

Shaban, A., Bansal, S., Liu, Z., Essa, I., & Boots, B. (2017). One-Shot Learning for Semantic Segmentation. Retrieved from http://arxiv.org/abs/1709.03410

Shankar, K. H., Singh, I., & Howard, M. W. (2016). Neural Mechanism to Simulate a Scale-Invariant Future. *Neural Computation*, *28*(12), 2594–2627. http://doi.org/10.1162/NECO_a_00891

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*(July 1928), 379–423. http://doi.org/10.1145/584091.584093

Smith, A. D., McKeith, L., & Howard, C. J. (2013). The development of path integration: Combining estimations of distance and heading. *Experimental Brain Research*, *231*(4), 445–455. http://doi.org/10.1007/s00221-013-3709-8

Smith, T. G., Lange, G. D., & Marks, W. B. (1996). Fractal methods and results in cellular morphology — dimensions, lacunarity and multifractals. *Journal of Neuroscience Methods*, *69*(2), 123–136. http://doi.org/10.1016/S0165-0270(96)00080-5

Smith, M. Lou, & Milner, B. (1981). The role of the right hippocampus in the recall of spatial location. *Neuropsychologia*, *19*(6), 781–793. http://doi.org/10.1016/0028-3932(81)90090-7

Spiers, H. J., Burgess, N., Hartley, T., Vargha-Khadem, F., & O'Keefe, J. (2001). Bilateral hippocampal pathology impairs topographical and episodic memory but not visual pattern matching. *Hippocampus*, *11*(6), 715–725. http://doi.org/10.1002/hipo.1087

Squire, L. R., & Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science (New York, N.Y.)*, *253*(5026), 1380–6. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1896849

Stern, C. E., Sherman, S. J., Kirchhoff, B. A., & Hasselmo, M. E. (2001). Medial temporal and prefrontal contributions to working memory tasks with novel and familiar stimuli. *Hippocampus*, *11*(4), 337–346. http://doi.org/10.1002/hipo.1048

Taube, J. S., Muller, R. U., & Ranck, J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *10*(2), 420–35. http://doi.org/10.1212/01.wnl.0000299117.48935.2e

Tavares, R. M., Mendelsohn, A., Grossman, Y., Williams, C. H., Shapiro, M., Trope, Y., & Schiller, D.

(2015). A Map for Social Navigation in the Human Brain. *Neuron*, *87*(1), 231–243. http://doi.org/10.1016/j.neuron.2015.06.011

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*(4), 189–208. http://doi.org/10.1037/h0061626

Trullier, O., Wiener, S. I., Berthoz, A., & Meyer, J.-A. (1997). Biologically Based Artificial Navigation Systems: Review and Prospects. *Progress in Neurobiology*, *51*(5), 483–544. http://doi.org/10.1016/S0301-0082(96)00060-3

Tulving, E. (1989). Memory: Performance, knowledge, and experience. *European Journal of Cognitive Psychology*. http://doi.org/10.1080/09541448908403069

Tulving, E. (2002). Episodic Memory: From Mind to Brain. *Annual Review of Psychology*, *53*(1), 1–25. http://doi.org/10.1146/annurev.psych.53.100901.135114

Umeyama, S. (1991). Least-Squares Estimation of Transformation Parameters Between Two Point Patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. http://doi.org/10.1109/34.88573

Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In *Analysis of visual behavior* (pp. 549–586). Cambridge: The MIT Press.

Uttal, D. H., & Chiong, C. (2004). Seeing space in more than one way: Children's use of higher order patterns in spatial memory and cognition. In G. L. Allen (Ed.), *Human spatial memory: Remembering where* (pp. 125–142). Mahwah, NJ: Erlbaum.

van der Spoel, E., Rozing, M. P., Houwing-Duistermaat, J. J., Eline Slagboom, P., Beekman, M., de Craen, A. J. M., … van Heemst, D. (2015). Siamese Neural Networks for One-Shot Image Recognition. *Aging*, *7*(11), 956–963. http://doi.org/10.1017/CBO9781107415324.004

Van Hoesen, G. W., Rosene, D. L., & Mesulam, M. M. (1979). Subicular input from temporal cortex in the rhesus monkey. *Science (New York, N.Y.)*, *205*(4406), 608–10. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/109926

Vargha-Khadem, F. (1997). Differential Effects of Early Hippocampal Pathology on Episodic and Semantic Memory. *Science*, *277*(5324), 376–380. http://doi.org/10.1126/science.277.5324.376

Voss, J. L., Bridge, D. J., Cohen, N. J., & Walker, J. A. (2017). A Closer Look at the Hippocampus and Memory. *Trends in Cognitive Sciences*, *xx*, 1–12. http://doi.org/10.1016/j.tics.2017.05.008

Wang, J. X., Cohen, N. J., & Voss, J. L. (2015). Covert rapid action-memory simulation (CRAMS): A hypothesis of hippocampal-prefrontal interactions for adaptive behavior. *Neurobiology of Learning and Memory*, *117*, 22–33. http://doi.org/10.1016/j.nlm.2014.04.003

Warren, D. E., Duff, M. C., Cohen, N. J., & Tranel, D. (2015). Hippocampus contributes to the maintenance but not the quality of visual information over time. *Learning & Memory*, *22*(1), 6–10. http://doi.org/10.1101/lm.037127.114

Warren, D. E., Duff, M. C., Jensen, U., Tranel, D., & Cohen, N. J. (2012). Hiding in plain view: Lesions of the medial temporal lobe impair online representation. *Hippocampus*, *22*(7), 1577–1588. http://doi.org/10.1002/hipo.21000

Watson, P. D., Voss, J. L., Warren, D. E., Tranel, D., & Cohen, N. J. (2013). Spatial reconstruction by

patients with hippocampal damage is dominated by relational memory errors. *Hippocampus*, *23*(7), 570–580. http://doi.org/10.1002/hipo.22115

Wilson, S. D. R., & Hulme, A. (1983). The Effect of Bubbles Attached to an Electrode on Electrical Resistance and Dissolved Gas Concentration. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *387*(1792), 133–146. http://doi.org/10.1098/rspa.1983.0054

Wolbers, T., & Hegarty, M. (2010). What determines our navigational abilities? *Trends in Cognitive Sciences*, *14*(3), 138–146. http://doi.org/10.1016/j.tics.2010.01.001

Wolf, M. J. P. (2010). Theorizing Navigable Space in Video Games. *Logic and Structure of the Computer Game*, 36–62.

Yee, T. S. L. (2012). *Medial Temporal Lobe and Prefrontal Cortex Contributions to Memory Expressed on Short Timescales*. University of Illinois Urbana-Champaign.

Yonelinas, A. P. (2013). The hippocampus supports high-resolution binding in the service of perception, working memory and long-term memory. *Behavioural Brain Research*, *254*, 34–44. http://doi.org/10.1016/j.bbr.2013.05.030

Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event perception: A mind-brain perspective. *Psychological Bulletin*, *133*(2), 273–293. http://doi.org/10.1037/0033-2909.133.2.273