3-5-2018 2:00 PM

Object processing in the medial temporal lobe: Influence of object domain

Anna Blumenthal, *The University of Western Ontario*

Supervisor: Köhler, Stefan, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Psychology

Recommended Citation

Blumenthal, Anna, "Object processing in the medial temporal lobe: Influence of object domain" (2018). *Electronic Thesis and Dissertation Repository*. 5311.
https://ir.lib.uwo.ca/etd/5311

Abstract

We live in a rich visual world, surrounded by many different kinds of objects. While we may not often reflect on it, our ability to recognize what an object is, detect whether an object is familiar or novel, and bring to mind our general knowledge about an object, are all essential components of adaptive behavior. In this dissertation, I investigate the neural basis of object representations, focusing on medial temporal lobe (MTL) structures, namely, perirhinal cortex, parahippocampal cortex, and hippocampus. I use what *type* of thing an object is, or more specifically, the broader category (e.g., “face” or “house”) or domain (e.g., “animate or “inanimate”) to which an object belongs to probe MTL structures. In the Chapter 2, I used fMRI to explore whether object representations in MTL structures were organized by animacy, and/or real-world size. I found domain-level organization in all three MTL structures, with a distinct pattern of domain organization in each structure. In Chapter 3, I examined whether recognition-memory signals for objects were organized by category and domain in the same MTL structures. I found no evidence of category or domain specificity in recognition memory-signals, but did reveal a distinction between novel and familiar object representations across all categories. Finally, in Chapter 4, I used a neuropsychological approach to discover a unique contribution of the hippocampus to object concepts. I found that an individual with developmental amnesia had normal intrinsic feature knowledge, but less extrinsic, or associative feature knowledge of concepts This decreased extrinsic feature knowledge led to abnormalities specific to non-living object concepts. These results show that the hippocampus may play an important role in the development of object concepts, potentially through the same relational binding mechanism that links objects and context in episodic memory. Taken together, these findings suggest that using object category or domain to probe the function of MTL structures is a useful approach for gaining a deeper understanding of the similarities and differences between MTL structures, and how they contribute more broadly to our perception and memory of the world.

Keywords

Medial Temporal Lobe, Perirhinal Cortex, Parahippocampal Cortex, Hippocampus, fMRI, Representational Similarity Analysis, Developmental Amnesia, Semantic Memory, Category-Specificity, Domain-Specificity, Animacy, Real-world Size, Object Concepts, Recognition-Memory

Co-Authorship Statement

The projects reported in the current thesis were carried out under the supervision of Dr. Stefan Köhler. The study presented in Chapter 2 is under revision for resubmission to *Human Brain Mapping*. Dr. Bobby Stojanoski is a co-first author on the paper, he contributed to design, data collection, and analyses. The paper also benefited from Dr. Chris Martin's theoretical and technical expertise, as well as Dr. Rhodri Cusack's technical expertise, reflected in their co-authorship. Chapter 3 is currently being prepared for publication with Jane Kouptsova, Dr. Martin, and Dr. Köhler. Dr. Martin and Jane contributed to data analysis. Lastly, the research presented in Chapter 4 is published in *Neuropsychologia* (Blumenthal et al., 2017). Dr. Devin Duke is a co-first author, he designed and implemented data collection, while analysis was done by him, myself, and Dr. Ben Bowles. Writing and revisions were done by myself, advised by Dr. Köhler and Dr. McRae. Dr. McRae's input at all stages of the project was crucial, reflected in his position as senior author. Additionally, Dr. Rosenbaum and Dr. Gilboa collaborated on this project, providing theoretical guidance. Lastly, all projects comprising this thesis were advanced through insights from fellow lab members, including Jane Kouptsova, Jordan DeKraker, HY Yang, and Kayla Ferko. The general introduction and discussion in this dissertation are largely my own ideas and writing, but were shaped by input from my advisor, Dr. Stefan Köhler.

Acknowledgments

First, I would like to thank my supervisor, Dr. Stefan Köhler. Stefan's mentorship over the last four years has been invaluable to my growth as a scientist. His passion for understanding the neural basis of human memory, and remarkably in depth knowledge of the literature has inspired me to pursue a career in this field. Stefan's meticulous approach to science and healthy skepticism has trained me to do my best work, and to always communicate my ideas precisely. Under his mentorship, I have become a stronger critical thinker, writer, speaker, and an engaged member of the cognitive neuroscience of memory community. I am grateful to have had this training opportunity, and look forward to continued collaboration!

I would also like to thank my advisory committee members: Dr. Jody Culham, Dr. Rhodri Cusack, and Dr. Ken McRae. Each provided valuable scientific and pragmatic advice. In addition to their official role on my committee, I had the opportunity to collaborate closely with both Dr.Cusack and Dr. McRae. These experiences truly made me understand the common claim that "no-one should ever do science alone." I learned an incredible amount from both of them, while always having a great time.

On the topic of not doing science alone, I am deeply grateful to all of the scientists at the BMI. The BMI has been a wonderful place to grow, thanks to the supporting, enthusiastic, and collaborative nature of its members. In particular, I would like to thank two BMI members who served as collaborators and mentors, Dr. Chris Martin and Dr. Bobby Stojanoski. Chris has been a role model for me, as his combination of theoretical knowledge and technical prowess are unmatched. His down to earth attitude has made him an excellent mentor and friend, always there to offer wise calmly-delivered advice (with a sprinkle of humor thrown in). Complimentary to Chris's even-toned advice throughout the years, has been Bobby's exuberant optimism, and endless energy for science. His support has kept me going through many challenges, and has shaped me into a better scientist and person. I am so grateful for his belief in me, and look forward to continued collaborations and mentorship.

I am also grateful to Dr. Jessica Grahn, for her career mentorship and encouragement of women in science more generally. I thank my lab members, in particular Jordan Dekraker, who is always willing to drop whatever he is doing to talk about the hippocampus, and Kayla Ferko, for being a fantastic mentee and the most entertaining data collection partner.

Finally, I would like to thank my family and friends for their support and enthusiasm throughout the years. Especially the unwavering support from those who didn't really understand why I was still in school, and in Canada nevertheless. My family has always understood and fostered my inner dork. They have encouraged me to pursue knowledge for the sake of knowledge, not let me complain too much, and are always there to remind me that studying the brain is pretty much the coolest thing ever. I couldn't have done this alone. Thanks everyone!

Table of Contents

List of Tables

List of Figures

List of Appendices

Chapter 1

1 Introduction

We live in a rich visual world, surrounded by many kinds of objects. While we may not often reflect on it, our ability to recognize what an object is, detect whether an object is familiar or novel, and bring to mind our general knowledge about an object, are essential components of adaptive behavior. Ultimately, these abilities guide how we interact with objects in our world. Our perception and interaction with objects is guided by what *type* of thing an object is, or more specifically, the broader category (e.g., “face” or “house”) or domain (e.g., “animate” like a face, or “inanimate” like a house) to which that object belongs. For example, imagine you are at the local park to meet your friends for a picnic. You are surrounded by many types of objects, including animate objects (dogs, people), small inanimate objects (your friend’s guitar, silverware for the picnic), and large inanimate objects (trees, park benches). When encountering an animate object, you may want to know if it is moving towards you or away from you, or whether it appears hostile and friendly, whereas with a small inanimate object you might want to pick it up, and use it to make music, or eat lunch. This is quite different from a large inanimate object, such as a tree, which you might use as a landmark to meet your friends for the picnic.

Interestingly, not only do object category and domain appear to guide our perception and behavior, but they are also reflected in our neural organization. One of the most striking examples is that brain damage can cause category or domain specific deficits (for a review see Capitani et al., 2003). Additionally, measuring activity in the healthy brain while participants view objects has revealed category and domain level organization across the ventral visual stream (VVS), a neural pathway involved in object processing running from occipital cortex through ventral to anterior lateral and medial temporal cortex (Mishkin & Ungerleider, 1982; Mishkin et al., 1983; Goodale & Milner, 1992; Goodale et al., 2008; Kravitz et al., 2013). Within the VVS, there are contiguous regions that show more activity for objects from a certain category than objects from other categories. Perhaps most well-known is the fusiform face area (FFA), a region on the right fusiform gyrus that shows higher neural activity for faces than for other object categories (Kanwisher et al., 1997), but similar regions have been reported for scenes (parahippocampal place area, Epstein et al., 1999), body parts (extrastriate body area, Downing et al., 2001), and words (visual word form area, McCandliss et al., 2003). Further, these category preferring regions are not randomly strewn throughout the ventral visual stream, but fall within a broader functional organization by domain. For example, the FFA falls within a larger zone of lateral cortex that shows more activity for animate objects than inanimate objects, whereas the PPA falls within a medial zone that prefers inanimate objects (for a review see Grill-Spector & Weiner, 2014). More recently, Konkle et al. (2013, 2012) have shown that this medial inanimate preference zone is for objects whose real-world size is large, and there is a separate zone located more dorsolaterally that shows a preference for small inanimate objects. This distinction by real-world size is not present in the animate domain, leaving this organizational principle to be referred to as the “tripartite distinction”. In addition to finding continuous regions of cortex that show more activation on average for one category or domain over another, similar category and domain level organization is observed when taking a less anatomically specific approach than looking for continuous preference zones, and examining the *representational space* of object-evoked responses across the VVS. Specifically, patterns of activity across the VVS are more densely clustered, or more similar to each other for certain categories of objects (e.g., faces), and category clustering is embedded within a broader domain organization, where animate objects are represented more similarly to each other than to inanimate objects, and vice versa (Kriegeskorte et al., 2008; Proklova et al., 2016). Real-world size also appears to matter in representational space, with large inanimate objects evoking more similar patterns of activity to each other than to small inanimate objects and vice versa (Julien et al., 2016).

Research on category and domain specificity in the VVS has focused on the more posterior and lateral portions of the temporal lobe, and typically has not included the full extent of the medial temporal lobe (MTL). MTL structures include perirhinal cortex (PrC), parahippocampal cortex (PhC), entorhinal cortex (EC), and the hippocampus (HpC). It should be noted that some studies have included a portion of the PhC, given the anterior aspect of the PPA is often functionally localized to both the lingual gyrus and the posterior portion of PhC. However, PhC in its entirety as defined anatomically has not been included in these studies. This lack of inclusion of the MTL in the vision literature is likely due to the fact that a prominent view has held that MTL structures are part of a system dedicated to declarative memory, and distinct from the VVS (Squire et al., 1991). However, evidence primarily from studies focused on the PrC has challenged this view. Specifically, evidence suggests that PrC is crucial for recognition memory of objects as well as fine-grained object perception (Murray and Bussey, 1999; Bussey et al., 2007). According to one theory, the representational hierarchical theory (R-H theory), PrC is crucial for both object memory and object perception because it represents the convergence of low-level visual features into holistic object representations (Murray and Richmond, 2001; Barense et al., 2012; Erez et al., 2015; for a review see Cowell et al., 2010).

The view that PrC is the apex of object processing in the VVS can also be seen in memory focused theories of the MTL. For example, according to the Binding of Items in Context model, or BIC model of episodic memory, PrC sends object information to the hippocampus where it is bound to context information to form an episodic memory (Eichenbaum et al., 2007; Diana et al., 2007 Davachi, 2006). In the BIC model, PhC represents context information, including spatial scene context. This notion aligns with work from the vision literature showing that the more posterior portion of PhC is often functionally localized by scenes (Epstein & Kanwisher, 1998). However, the picture becomes more complicated when considering that posterior PhC also shows a preference for buildings and other large objects (Aguirre et al., 1998; Epstein & Kanwisher, 1998; Bar & Aminoff, 2003; Mullally & Maguire, 2011; Konkle & Oliva, 2012; Troiani et al., 2012; Magri et al., 2016). Further, Martin et al. (2013, 2016) has shown category-specific recognition memory signals for objects in PhC. Thus, it remains an open question as to whether PrC is involved in object processing in a domain-general manner, or whether both PrC and PhC contribute to object processing, but differ by the category or domains to which they are sensitive. The HpC is generally not considered as being involved in object processing, or to show category or domain level organization, but this has not been thoroughly tested. Specifically, the HpC is thought to receive both object and spatial context information, and bind that information together, resulting in representations that are a combination of objects from different categories and their spatial context - and are therefore category and domain general (for a review on the BIC model see Eichenbaum et al., 2007).

In the current dissertation, I used object domain as a tool to explore how MTL structures contribute to object perception, memory, and semantic knowledge of objects. To investigate these questions, I used two approaches, fMRI in healthy adults, and behavioral testing in individuals with selective damage to MTL structures. Specifically, in Chapter 2, I explore how object representations are organized by category and domain across MTL structures using fMRI. In Chapter 3, I interrogate the same fMRI data to ask whether object *recognition-memory* signals themselves are organized by category or domain in MTL structures, and whether memory is coded by average changes in signal, or by pattern based changes in representation. In Chapter 4, I explore a unique way that one MTL structure, the HpC, may contribute to object knowledge specific to the nonliving domain.

To motivate the following chapters, in the remainder of the Introduction I first provide a brief overview of the evidence for category and domain specific neural organization in three parts: 1) category and domain specific deficits in neuropsychological research; 2) fMRI evidence for category and domain specific VVS organization in anatomical space; and 3) fMRI evidence for category and domain specific VVS organization in representational space. I then briefly cover R-H theory because it provides motivation for viewing MTL regions as a continuous part of the VVS, and therefore for exploring organization in a similar manner. Additionally, I introduce how R-H theory models of PrC motivated us to investigate whether memory signaling in this region was category or domain specific. I then provide a brief overview of the BIC model, which further motivated exploring content differences across MTL structures, and whose HpC model formed the basis for our investigation of how the HpC might contribute to semantic knowledge of objects. I then discuss current fMRI work on content differences, in particular category specificity, across different MTL structures. Finally, I provide a brief overview of feature-based models of object concepts, and how they differ by domain.

1.1 Category and domain specific deficits in neuropsychological research

One well-known type of category-specific impairment is prosopagnosia, or “face blindness”, which is a neurological disorder in which individuals are impaired specifically at recognizing faces. Prosopagnosia can be acute (acquired) or congenital (Behrmann & Avidal, 2005). This disorder generally occurs when there is damage to the fusiform face area (Gruter et al., 2008), or to connections from the FFA to other brain areas (Thomas et al., 2009). Selective impairments have also been reported for knowledge of tools (see Johnson-Fray, 2004 for a review). While to our knowledge there are no cases of specific impairments for buildings, cases of landmark agnosia have been reported, in which individuals can no longer recognize or orient themselves in their environment by familiar landmarks, generally after damage to the PPA (Takahashi et al., 2002; Claessen et al., 2017). Overall, these results suggest that there is some category specificity in the neural organization of object knowledge, at least for some categories.

At the broader domain level, a number of studies have described individuals with domain specific deficits in object knowledge due to focal brain damage (for a review see Capitani et al., 2003). The earliest report of disproportionate impairments in certain categories of knowledge over others were those by Warrington, Shallice, and McCarthy (1984). They reported 4 patients who had made a partial recovery from herpes encephalitis, all of whom showed extensive bilateral temporal damage on CT scans. Interestingly, each patient showed significantly more impairment for knowledge of living things and foods than for nonliving objects. These impairments occurred in both the verbal and visual domain. This pattern of more severe impairment for living objects relative to inanimate holds in the majority of cases. Capitani et al., (2003) reviewed 79 published case studies of category specific deficits, and found 61 cases with individuals who had a disproportionate impairment for biological categories relative to artefacts, and 18 showing a disproportionate impairment for artefacts. Further, they argue that patterns of category selective impairment can be fractionated into animate objects, inanimate biological objects (fruits and vegetables), and artefacts. While there are some cases of more highly specific category impairments, most commonly the impairment covers a broader domain, with an obvious distinction between living things and artefacts. What about evidence for a distinction between large and small inanimate objects? This is somewhat less clear in the current literature, as usually objects used in testing are smaller (tools) so these dimensions have not been thoroughly pit against each other in the neuropsychology literature.

1.2 Category and domain organization: Anatomical space

One of the most well-known examples of category-specific neural organization was the discovery of the fusiform face area (FFA). Kanwisher et al. (1997) showed participants images of faces as well as images of other object types, and found that a univariate contrast to evaluate whether any regions show higher average activity for faces revealed a contiguous region in the right fusiform gyrus appears fairly robustly across individuals. This finding has since been widely replicated, and has sparked a lively debate as to whether the FFA is a module dedicated to face processing (Grill-Spector et al, 2004; Grill-Spector et al., 2006; Gauthier et al., 1999, 2000). Importantly, it is possible to decode patterns of activity specific to other object categories within the FFA (Haxby, 2001), suggesting that it may not be exclusively involved in processing faces. However, a region that prefers faces, even if not dedicated, does suggest some category-level functional organization of visual cortex. Additionally, a number of other regions have been found in occipital and temporal cortex that show differential increase in activity for specific categories, including, for example, for scenes, places, and buildings (parahippocampal place area, PPA) (Epstein & Kanwisher, 1998; Epstein et al., 1999), for bodies (extrastriate body area) (Downing et al., 2001), and for tools (Chao et al., 1999).

Interestingly, regions with preferential responses to specific categories are not organized randomly in their spatial relationships. Instead, they fall within broader functional zones, where inanimate objects more strongly activate a large swath of cortex medial to the mid fusiform sulcus, and animate objects more strongly activate the area lateral to the sulcus (for a review see Grill-Spector and Weiner, 2014). In other words, the FFA and EBA are embedded within a larger animacy zone, and the PPA within a larger inanimate zone. More recently, fMRI based work by Konkle et al. (2012, 2013) has shown that this inanimate medial zone is specialized for inanimate objects whose real-world size is large, whereas there is a separate zone located more dorsal and lateral to the animate zone that prefers objects whose real-world size is small (including a tool-preferring region). There is no evidence for an equivalent divide in zones by real-world size for animate objects, leaving this organization schema to be referred to as “tripartite” (Konkle et al., 2013). Importantly, these domain preference zones can be identified in a variety of tasks, whether it be passive viewing of images, low level perceptual tasks, or object categorization, but the majority of work in this area has used fairly low-level perceptual tasks to ensure subjects are paying attention. Some work has shown that task might impact the involvement of more posterior or anterior VVS regions, for example, Taylor et al. (2012) show more posterior areas are involved in naming object domain, whereas naming a specific object engages more anterior areas. It is an open question therefore, as to whether anterior areas such as the MTL which may be more involved in processing individual objects, still show category and domain organization.

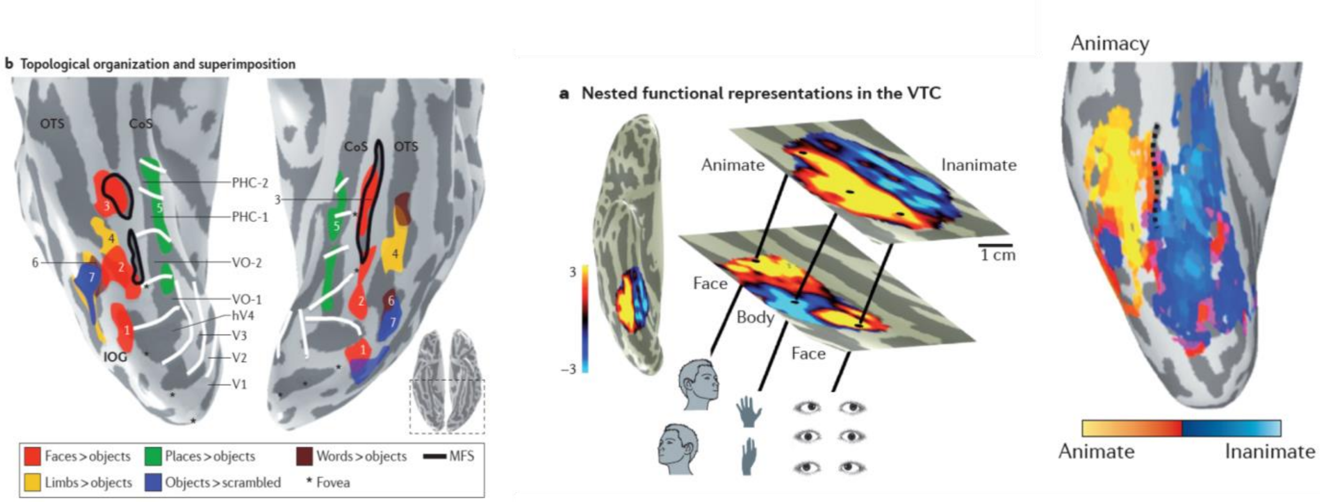

Figure 1.1: Category and domain level anatomical organization in the ventral visual stream (Grill-Spector & Weiner, 2014).

1.3 Category and domain organization: Representational space

Evidence from univariate analyses, as described above, has shown contiguous regions in the VVS that show more activity for certain categories or domains of objects. In this approach, activity is averaged across voxels for each category or domain, and then average activity levels are contrasted between categories or domains, and thresholded by looking for continuous clusters of voxels that show a similar response. Both category and domain level organization of visual cortex can also be seen using a quite different analysis approach, representational similarity analysis (RSA; for a review see Kriegeskorte et al., 2008b). In RSA, brain activity is measured by assessing how fine-grained patterns of activity evoked by different objects relate to one another. RSA, therefore, does not require that voxels in the same areas across subjects show the same response patterns, as stimuli are now compared in “representational space”, where a “representation” is defined by the pattern of activity evoked across voxels in a given region in response to a distinct stimuli. For example, Kriegeskorte et al. (2008a) showed participants a large number of objects spanning a number of categories and domains, while participants performed a low-level visual task (press a button when the fixation cross appears red). Patterns of activity evoked across the entire VVS for each object were compared to the patterns of activity evoked for each of the other objects using a dissimilarity measurement (1-Pearson’s correlation), and were then plotted in a representational dissimilarity matrix (RDM) (see fig. 3). From this data driven approach one can see both category and domain level organization of object evoked responses. For example, there is a dark blue square indicating the responses for face stimuli compared to other face stimuli are more highly similar to each other than they are to other object categories. Perhaps even more striking is the domain level organization across all stimuli, in which animate objects are more similar to each other than they are to inanimate objects, and vice versa. Interestingly, a similar representational space can be revealed in nonhuman primates when comparing patterns of activity in neural firing across visual cortex during viewing of the same stimuli (Kriegeskorte et al., 2008a).

Using a similar RSA approach, a recent study showed that real-world size is an organizing principle across representational space in visual cortex (Julien et al., 2016). Participants were scanned while they viewed inanimate objects from 20 different categories, whose real world size was large or small (while retinal size was controlled for). They were instructed to memorize each object to ensure that they paid careful attention. Patterns of activity evoked by large objects were more similar to other large objects than they were to small objects and vice versa, when evaluated across a large region spanning temporo-parietal-occipital cortex, similar to the large cortical volume examined by Kriegeskorte et al. (2008). This pattern of organization was also apparent when evaluating representations in smaller functionally defined regions known to show average increased activity for scenes or large objects (PPA, retrosplenial cortex, occipital place area), as well as object responsive regions that do not show increased average activity for large objects (lateral occipital cortex, posterior fusiform gyrus), but this organization was not present in early visual cortex.

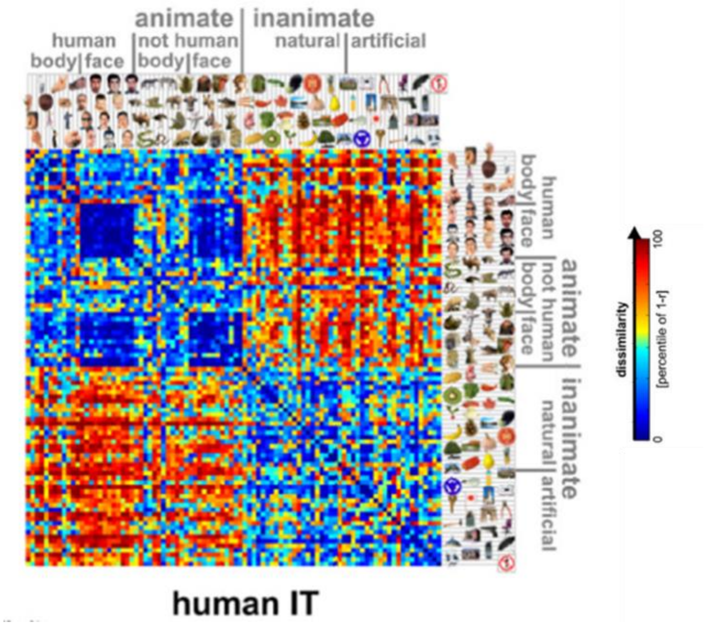

Figure 1.2: Category and domain level organization of representational space in the ventral visual stream (Kriegkeskorte et al., 2008).

1.4 Representational-Hierarchical theory

Research on category and domain-level functional organization of object evoked responses in the brain has centered on the occipital lobe, and on the posterior and lateral temporal lobe, generally leaving out the anterior medial temporal lobe, which includes PhC, PrC, EC, and the HpC. The posterior portion of PhC is sometimes included in these studies because the anterior portion of the PPA is localized here. However, the entire structure defined anatomically, as well as the PrC, EC and HpC, are typically not examined. Some of the reasons for this might be historical, as MTL structures were generally considered to be part of a dedicated declarative memory system (Squire et al., 1991), and therefore outside of the purview of object recognition. Indeed, category and domain organization has been mostly studied in passive viewing paradigms or during low-level perceptual tasks. However, at least one theory, the representational-hierarchical theory (R-H theory) has challenged the view that MTL structures should be considered as dedicated to memory, and argued for the inclusion of these structures as an extension of the VVS (Murray & Bussey, 1999; Murray & Richmond, 2001; Saksida & Bussey, 2010; Cowell et al., 2010). The R-H theory has guided this research by motivating us to extend category and domain specific mapping to MTL structures, and to ask questions about links between memory signaling and representational content. The central tenet of R-H theory is that, in general, there are no brain regions dedicated to one specific psychologically defined *process*, such as memory or perception. Instead, functional differences between brain regions are better characterized by their representational content, or the form of information they carry. This content may be used in different processes, e.g., perception or memory. In what follows, I will describe evidence for the R-H view of MTL structures.

According to R-H theory, the key difference across VVS regions, from early visual cortex through HpC, is that representations become more highly conjunctive as you move along the posterior to anterior axis of occipital and ventral temporal cortex. Therefore, brain damage to a particular region should disrupt performance on any task for which representations at that level of complexity and specificity are required. For example, in pairwise visual discrimination tasks of objects, PrC lesions in non-human primates and rodents do not impair performance when the objects can be discriminated by a single low-level feature, but impair performance when feature conjunctions must be used (i.e., when feature ambiguity is high, or individual features are highly overlapping) (Buckley and Gaffan, 1998; Bussey et al., 2002, 2003; Bartko et al., 2007). The same pattern is seen in the HpC, for even more highly conjunctive content than single objects, for example, in the discrimination of complex scenes with spatial object conjunctions (Lee et al., 2006; 2012). Overall, a role of the MTL in visual perception has been reported in rodents, non-human primates, and humans, using a variety of methodologies, providing strong evidence for a causal role of MTL structures in perception, as opposed to just memory, although there is still an avid debate (for reviews see, Graham et al., 2010; Saksida and Bussey, 2010; Murray et al., 2007; but see Squire & Zola-Morgan, 1991; Squire et al., 2004; Suzuki et al., 2004, Suzuki, 2009a, 2009b).

Research guided by the R-H view has also accounted for patterns of memory impairment in cases of MTL damage. One powerful approach has been to test predictions derived from computational models based solely on a network organized by conjunctive representations. Cowell et al. (2006) created a connectionist model in which inputs were object features, and a lower layer represented caudal, or more posterior VVS areas with feature-based representations, and an upper layer corresponding the PrC, which represented the conjunctions of those features (whole objects). When the PrC layer of the model was lesioned, susceptibility to interference was increased, as there was no way to disambiguate objects with overlapping features relying solely on the caudal layer. The same susceptibility to interference was shown in rats with PrC lesions (Bartko et al., 2010, but see Clark et al., 2011), and humans with amnesia (Barense et al., 2012, but see Kim et al., 2011; Suzuki, 2009). Similarly, increased impairments in object recognition memory are seen with delays, as a result of feature interference from stimuli encountered during the delay period. This feature interference from stimuli is thought to be akin to real-life situations, during which there is almost always a continuous stream of visual information incoming. This computational evidence provides an explanation for delay-dependent impairments on discrimination tasks in human and non-human animals (Meunier et al., 1993; Mumby and Pinel, 1994). This model also generated a new prediction - which impairment in amnesia should actually be due to detecting novel stimuli as familiar, because without high-level conjunctive representations, the high amount of feature overlap in novel stimuli should lead to familiarity signals. McTighe et al. (2010) tested this prediction by familiarizing rodents with an object during a study period, and then by putting rodents in a dark cage (no visual interference), or a regular cage (visual interference) for a delay period. Rodents were tested to determine how they explored novel or familiar stimuli (rodents spend more time exploring novel stimuli naturally). Rodents in the visual interference condition treated novel stimuli as familiar, but this pattern of results was not present in the group with no visual interference. These results suggest that delay-dependent memory impairments can be accounted for by considering representations in PrC alone. Without high-level object representations, there is a high amount of feature interference, which causes novel objects to appear familiar. Importantly, this changes the mainstream conception of amnesia, as an inability to recognize things as familiar, to a discrimination problem where *everything* feels familiar.

These findings provide strong evidence for R-H theory, and have motivated consideration of MTL structures as a part of the VVS in the current thesis. If one considers MTL structures to be a part of the VVS, it is therefore important to explore the organization of object representations in these regions along the same category and domain lines that have previously been explored in more posterior and lateral VVS regions. Further, computational R-H models provide a mechanism by which representations in these regions can lead to memory signals, thereby directly motivating our analyses in Chapter 3, where we asked if memory signals were organized by category and domain. Specifically, connectionist models represent objects with overlapping features more similarly, and recognition memory signals come directly from activation of these representations (Cowell et al., 2006). Finally, this view encourages “thinking out of the box” in terms of how MTL representations might contribute to any process that requires a particular form of content, closely in line with our finding in the final chapter that HpC contributes to semantic memory (and not only episodic memory) through creating representations of items linked to contextual information.

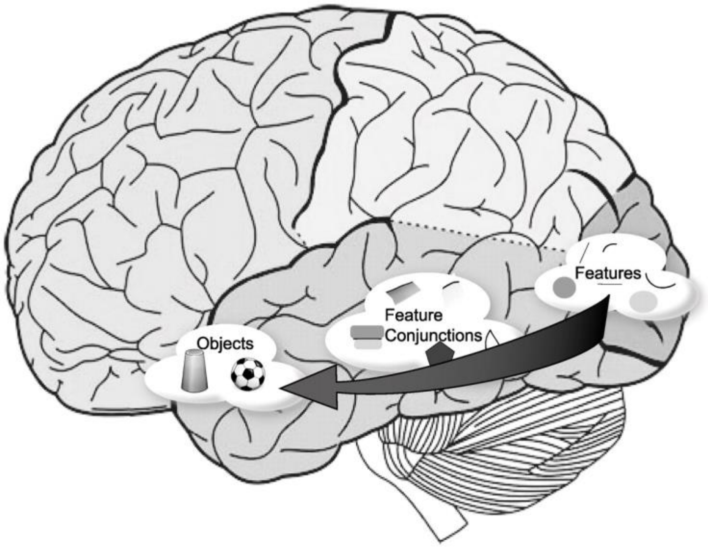

Figure 1.5. Representational-Hierarchical theory (Barense et al., 2012)

1.5 Binding of items in context model

While R-H theory has focused on an inclusion of MTL structures as an extension of the VVS, with an emphasis on content, characterizations of MTL structures with a focus on their contribution to long term memory share some overlap with R-H theory. In particular, the binding of items in context model, or BIC model, has focused on how MTL structures differ in their contributions to long term memory in terms of the types of information each structure contributes. Like the R-H model, the BIC model also places PrC at the apex of the VVS, containing high level object representations. Similar to much of the empirical work on object processing in PrC, the model treats PrC as being domain general, leaving an investigation of domain and category specificity unexamined. In terms of the PhC, the BIC model argues that this structure represents spatial context, a view that must be reconciled with its role in object specific processing, which work in our lab and in the current dissertation addresses. Finally, the BIC model has focused on the role of the HpC as binding item and context information for the service of episodic memory, a role which we expand in the current thesis to domain specific object knowledge.

The BIC model was proposed Eichenbaum et al. (2007), although a similar model (without the addition of functional differentiation between lateral and medial entorhinal cortex) was proposed by Davachi (2006). The model describes two pathways converging in the HpC, the VVS pathway that culminates in the PrC and lateral entorhinal cortex, which contains high-level item information, and a dorsal stream pathway that culminates in the PhC and medial entorhinal cortex containing context, including “where” information. Because both the lateral and medial entorhinal cortex provide input to the HpC, these two streams converge in the HpC. It is thought that a key function of the HpC is to bind this item and context information into a holistic representation – the basis for rich coherent episodic memory (Eichenbaum et al., 2007; Davachi, 2006; Diana et al., 2007, first proposed by Mishkin, 1983). Importantly, the model includes the assumption that information flow is not purely hierarchical; there are projections from the hippocampus through the EC back to PrC and PhC, which have their own back projections to more posterior visual areas. It should be noted that emphases on the lateral and medial EC initially came from animal work, given EC is difficult to image in humans, although recent work has found initial evidence for a functional distinction in humans (Maass et al., 2015). According to this model, the functional role of back-projections is to reinstate previously encountered information, providing the neural basis for episodic recollection. For example, a single cue, such as an object, can activate a previously stored HpC pattern, which can then cause reinstatement of an episode by reactivating the original object and context information in the ventral and dorsal visual stream. In what follows, I overview some evidence for the BIC model of MTL function, and then discuss some of the shortcomings of the model. Finally, I discuss this model in relation to the R-H model and the questions about MTL structures posed in this dissertation.

Evidence supporting the BIC model comes from fMRI, human MTL lesion studies, as well as non-human animal models. Some of the evidence arose from attempts to understand whether MTL regions contribute differentially to two subjective experiences of remembering, termed familiarity and recollection. Familiarity refers to the experience of feeling a sense of familiarity for an object or entity (a face you see on the bus), while recollection refers to recalling contextual information specific to your initial encounter with that object or entity (recalling that the person on the bus is a butcher, and a time you went to buy a pork-chop in the shop last week). Notably, the content of these two experiences differ, in the case of familiarity you have only an item or object, and in the case of recollection you re-experience an episode with that object/entity, but also contextual information (and often multiple other objects or entities). These two types of remembering involve different MTL regions, patients with selective HpC damage have more difficulty with recollection, but preserved item familiarity (Yonelinas et al., 2002; Aggleton et al., 2005), and a patient with PrC/EC damage and a preserved HpC showed abnormalities in familiarity but not recollection (Bowles et al., 2007), although it should be noted that the damage also included anterior lateral temporal cortex. Interestingly, while most evidence suggests a dichotomy between HpC and PrC, one patient with damage more posterior, in PhC, also showed more impairments in recollection than familiarity (Cipolotti et al., 2006). Similar dissociations are seen in non-human animals. For example, monkeys and rodents with PrC lesions can no longer perform delayed-non-match-to-sample tasks in which they are required to identify which object is novel after a delay period, whereas those with damage restricted to the HpC still can (Nemanic et al., 2004; Mumby, 2001). PhC and HpC do impair object recognition, but only when the novelty consists of putting the object in a new location or context (Eacott and Norman, 2004; Mumby et al., 2002). Lastly, electrophysiology shows neurons in PrC are sensitive to changes in the amount of experience with an object, whereas HpC and PhC neurons do not show the same selectivity (but show stronger spatial coding) (for a review see Brown and Xiang, 1998).

In terms of fMRI studies, there is also evidence to suggest that PrC is involved primarily in object/item processing, PhC context, and the HpC in the binding of items in content. The evidence is both rich and complex, and here I provide only a brief overview of some examples. At encoding, HpC activity is higher for items that are later recollected, and is also higher during recollection (for review see Eichenbaum et al., 2007). PrC, on the other hand, has been shown to track familiarity at encoding, by showing more activity for items later associated with a stronger familiarity response (Ranganath et al., 2004), and has also shown decreased responding at retrieval that tracks the strength of subjective item familiarity (Yassa and Stark, 2008). This PrC BOLD pattern of response is similar to the pattern of responses in neurons recorded in monkey and rodent PrC. HpC activity has been found when participants remember whether two items were presented together (Kirwan et al., 2004; Jackson et al., 2004; Sperling et al., 2003), remembering the context an item was presented in (Davachi et al., 2003; Ranganath et al., 2004), or remembering the location an item was presented in (Staresina et al., 2006). A key difference seems to be whether the encoding of information is within the object, or associative by nature. For example, Diana et al. (2009) had subjects either encode color as part of an object “The elephant is red because it is sunburned” or external to the object “The elephant stopped at the light because it was red”. At retrieval, HpC and PhC activity correlated with recalling the associative encoding, whereas PrC activity correlated with successful retrieval only when the color was encoded as a feature of the object.

Unlike the R-H model, the BIC model is framed specifically around explaining how MTL structure contribute to different types of memory. At first glance this focus on the division between PrC and HpC in terms of their contributions to recollection or familiarity may seem antithetical to an R-H view, given that the R-H view stresses that brain structures are not modules dedicated to processes. However, a link can be made when considering how the representations in PrC and HpC might differently contribute to the subjective states of recollection and familiarity. Specifically, representations of objects/entities in PrC contribute to familiarity, which is by definition based on a signal from an object or entity alone. On the other hand, associative or items-in-context representations in HpC necessarily contribute to recollection because it is defined by the experience of this type of information (see Graham et al., 2010; Cowell et al., 2010). Indeed, the idea that moving beyond process memory based distinctions to evaluating MTL structures based on content or representations can be clearly understood from the BIC model perspective. For example, Davachi (2006) concludes by emphasizing that

"Future work should focus on revealing the nature of representational capacities of the MTL cortical input structures" (Davachi, 2006, pg. 698).

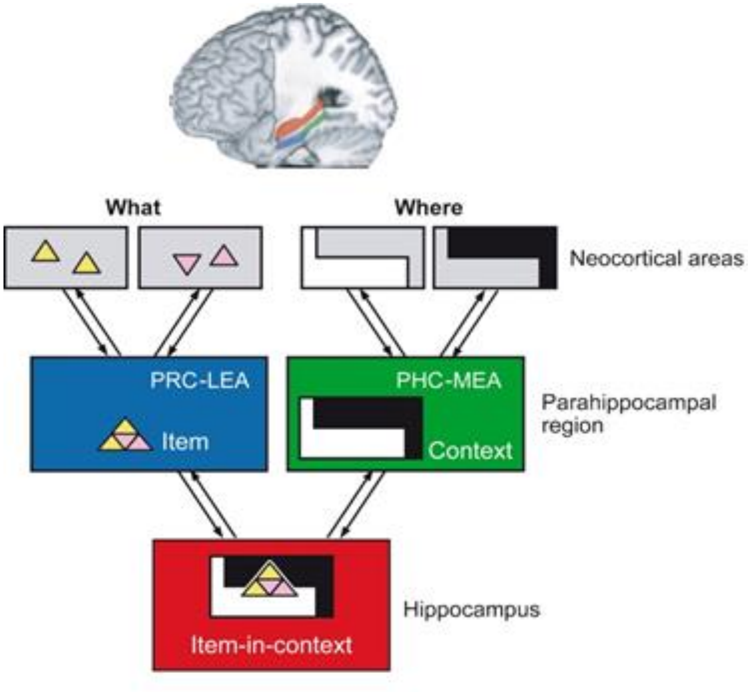

Figure 1.6. Binding of items-in-context (BIC) model (Eichenbaum et al., 2007).

1.6 Object processing in the medial temporal lobes

In light of characterizing MTL structures based on their differential content, a number of questions still remain, one of which pertains to PhC content. In the BIC model, PhC has been characterized as being involved in context (i.e., remembering a scene that is associated with a face), and spatial information more generally. One obvious challenge to this characterization of PhC is that in the vision literature the PPA (which partially covers posterior PhC) shows sensitivity to buildings and other large objects (Magri et al., *under review*). Martin et al. (2013, 2017) have shown recognition memory signals for objects from some categories in PhC, namely buildings, furniture, and trees. However, aside from this research, there have not been many other experiments exploring PhC object-based content, particularly during object memory tasks. Therefore, it remains an open question as to whether PrC represents most types of objects, or how object content is divided across regions. While some studies suggest interesting content differences (see next section), this has not been mapped out with a large number of object categories or by object domain.

A number of studies have explored differences in content, or representations, in MTL structures. For the most part, these studies have used stimuli from a small number of categories shown to be prominent in VVS organization more posteriorly, such as faces or scenes, and generally treated objects as a single undifferentiated category. However, even with this small sample of stimuli types, interesting patterns of category selectivity emerge across structures. For example, Litman et al. (2010) had participants perform a 1-back on images of scenes, faces, and objects. They found that there was a shift in average activity from PhC to PrC along a posterior to anterior gradient, with the strongest responses to scenes present in the most posterior part of the parahippocampal gyrus, and objects and faces in more anterior regions. While there was a shift in preference, it is notable that both structures showed above baseline responses to all stimuli. Similarly, in a target detection task, Liang et al (2012) found increased responding for scenes in PhC and faces in PrC. Neither study found any stimuli type preference in HpC response.

Multivariate analyses, exploring patterns of activity evoked in response to different stimuli types across MTL structures, have found somewhat similar results. Liang et al. (2012) decoded both scene and face stimuli from PhC and PrC, but decoding accuracy was higher for scenes in PhC than PrC. Similar results were found when participants viewed images and completed a low level perceptual task (indicating when a border came up on the image), or judged the image as pleasant or unpleasant. Decoding scenes from faces was highest in PhC, but objects versus faces could also be decoded above chance (Huffman and Stark, 2014). In PrC, the highest decoding accuracy was between faces and objects, but was also above chance between faces and scenes. Diana et al. (2008) used an expanded stimuli set in a 1-back task, which included scenes, faces, toys, objects, and abstract shapes. In PhC scenes were had the highest decoding accuracy, indicating the most distinct categorical responses relative to other object categories. However, it was also possible to distinguish faces and toys from other categories, which the authors interpreted as evidence that PhC does not solely represent scene information. LaRocque et al. (2013) explored patterns of responses in MTL structures while participants performed a low-level perceptual task (i.e., to press a button when fixation changes colors) on objects from many different categories, scenes, and faces. They found distinct object representations in both PhC and PrC, scenes in PhC and faces in PrC. In the majority of these studies, the HpC showed no stimulus category selectivity, gaining it the title of being “agnostic” (Huffman and Stark, 2014), although in one study scene information was decoded in the posterior portion of hippocampus (Liang et al., 2012). This evidence supports the idea that the HpC is agnostic because it is a convergence zone for both object and spatial information. This information is thought to be bound together to form a distinct episode, that would be category and domain general because it is a unique conjunction of different kinds of information. Further, it is thought that a key function of the HpC is pattern separation, or the process of orthogonalizing representations to reduce interference in memory (for a review see Yassa & Stark, 2003), a process that would also reduce any shared category information between stimuli.

While the studies outlined above have evaluated whether responses of MTL structures to stimuli are organized by category based on activity evoked while the stimuli are displayed and participants are performing a fairly low level task, to our knowledge only one set of studies has evaluated whether recognition memory signals in these regions (i.e., patterns of response to perceived novel or familiar stimuli) are categorical in nature, and whether this differs across structures. Specifically, Martin et al. (2013, 2017) asked for which categories it was possible to decode the distinction between perceived novel and familiar stimuli in each MTL structure. Importantly, in these studies any contribution of PhC due to context was removed by focusing exclusively on item-based familiarity responses. In PrC it was possible to decode memory signals for faces, furniture, and planes, and in PhC buildings, furniture and trees. The authors argue that a key dimension guiding PhC sensitivity to object category is navigational relevance, given that buildings, furniture, and trees are all stable and can be used for navigation, whereas planes, are highly mobile and therefore are less suited to navigation. In these studies it was not possible to decode category specific memory signals for any category in the hippocampus.

While these findings point to interesting differences and similarities in content representation between PhC and PrC, and a likely agnostic HpC, the small number of stimuli types used and/or a lack of differentiating object types makes it hard to understand more broadly what the differences are between structures. Further, the relation between category specificity in response to stimuli, or in memory coding, or both, is still unclear. Both of these gaps are addressed in Chapter 2 and Chapter 3 of this dissertation.

1.7 Feature-based object concept models

Thus far, we have considered the importance of object recognition (identifying an object) and object recognition memory (identifying an object as novel of familiar), and the neural organization of object responses by category and domain across VVS and MTL. However, we have yet to discuss how any such organization relates to stored concepts (i.e., object knowledge) derived from years of experience. Here I use the term “object concept” to refer to our semantic knowledge of an object. For example, our concept of a hammer includes what it looks like (wooden handle, metal head) but also what its function is, how we use it, and where we might find it. Indeed, thinking about what our knowledge of objects entails, and how that might drive brain organization, has a rich literature (Mahon & Caramazza, 2009; Martin et al., 2000; Martin et al., 2001; McRae & Cree, 2002). This literature can be used to derive hypotheses about how MTL structures contribute to object knowledge, and how it might differ by domain or category. Here, I cover only a slim portion of this vast literature, focusing on feature-based models of object concepts. First I cover how feature-based models have been used to explore the role of MTL structures, in particular PrC, in terms of how it represents objects, and how this is linked to object perception and memory. Second, while most attention to feature-based object models in the MTL has been in using them to understand PrC function, we show that they can be expanded to explore a previously unknown contribution of the HpC to object processing. This motivated the study presented in Chapter 4, where we investigated whether the HpC is involved in conceptual representations, and whether this involvement differs by object domain.

According to feature-based models of object concepts, object concepts are composed of a number of features bound together. Characterizing objects based on the relations among their features is a useful approach for understanding how they relate to each other in terms of category and domain organization in both psychological and brain space. Specifically, one can build a model of how objects relate to each other in semantic space, in terms of how many features they share (Cree and McRae, 2003). Semantic space can be derived from normative studies in which participants are presented with a concrete concept and list as many features as they can think of that make up that concept. Models based on these features have been shown to reflect representational space in PrC when participants name visually presented objects (Clark & Tyler, 2014) and perform a property verification task on object words (specific to left PrC; Bruffaerts et al., 2013). Importantly, feature-based object models account for patterns of semantic impairments after brain damage, such as the observation that impairment in knowledge for living things occurs more often than impairments for nonliving things (Cree and McRae, 2003). This distinction between living and nonliving things can be captured by differences in feature statistics, in particular living things are more highly similar to each other in terms of semantic feature overlap (Taylor et al., 2012). Tyler et al. (2013) used a precise feature-based statistical measure, correlation by distinctiveness, to capture the challenge of differentiating similar objects, and showed that this measure modulates bold activity in left PrC during a picture naming task (Tyler et al., 2013). Overlap in semantic features has also been shown to be causally related to PrC function. Specifically, Kivisaari et al. (2012) examined the volume of PrC, EC, and HpC in individuals with varying levels of atrophy due to Alzheimer’s disease. They found that volume in PrC, but not the other structures, predicted latencies when naming living, but not nonliving things, they concluded that these results show the importance of PrC in disambiguating semantically similar objects. Indeed, PrC has been shown to show more average activity when naming living than nonliving things (Bruffaerts et al., 2013).

Beyond differences between the living and nonliving domain by feature distinctiveness and feature overlap, these domains differ in terms of the types of features that are salient to their representations. For example, features can be classified into knowledge types, which can be further subdivided by whether those types of knowledge are *intrinsic* to, or a part of the object itself (i.e., color), or *extrinsic* to the object (i.e., function of the object, or location it is found; McRae et al., 2005). The importance of intrinsic versus extrinsic feature knowledge has been shown to differ by whether an object is living or nonliving. Specifically, extrinsic knowledge is more important for nonliving objects, whereas intrinsic knowledge is more central for living things (Barr and Kaplan, 1987). While the terminology differs in the semantics literature, we note that extrinsic features could be considered "associative" or "contextual" in nature. This opens up the possibility that while the HpC is usually thought not to be involved in object processing, and to be "agnostic" in terms of object category and domain information, it may play a previously undiscovered role in people's extrinsic semantic knowledge, and therefore in representing one domain of objects - namely nonliving objects.

1.8 Summary of literature review

In summary, object processing, whether it be recognition of an object, detecting whether that object is familiar or novel, or using semantic object knowledge, is a crucial part of human behavior. Unsurprisingly then, a large swath of the cortex, in particular the occipital and temporal lobe, are involved in object processing through a hierarchical information stream - the ventral visual stream. This neural basis of object processing is organized, both anatomically, as seen in category preference regions and larger domain preferring regions (i.e., the tripartite organizing schema), and in representational space. However, anterior medial temporal lobe regions have been less thoroughly explored, despite the fact that content may play a key role in their differential involvement in both object perception and memory, according to at least two key theories- the R-H theory and the BIC model. While there is some evidence of differential organization by stimuli type, with PhC being sensitive to scenes and large objects, and PrC perhaps more involved in faces, a careful examination across a large number of object categories is still needed. Further, examining domain organization in these regions has not been done, and can be a useful way to characterize their differential contributions. Additionally, a further examination of the link between stimuli based categorical responses and their relation to object memory signaling in these regions can potentially provide support for an R-H based account of MTL function. Lastly, we can use insights from the object concept literature to further test the limits of how MTL regions are involved in object processing for certain domains.

1.9 Goals of the current dissertation

In this dissertation, I address these gaps through three empirical studies. In Chapter 2, I first explore object category and domain organization for a large number of categories in each MTL structure, in the context of a continuous recognition memory task for objects. In Chapter 3, I ask whether the memory signals for these objects themselves are organized by category and domain, and whether this differs across structures in relation to their stimuli-response based organization. Further, I ask whether memory status is coded by repetition suppression, or pattern based representational changes and relate this to the R-H model that suggests memory is computed directly from stimuli representations. In Chapter 4, I turn to one specific MTL region, the HpC, often thought not to be involved in object processing and agnostic to domain, and provide some evidence that challenges these views. The important implications of this research are tied to the idea that the HpC binds items-in-context, or is involved in representations that have object and spatiotemporal conjunctivity.

1.10 References

Aggleton, J. P., Vann, S. D., Denby, C., Dix, S., Mayes, A. R., Roberts, N., & Yonelinas, A. P. (2005). Sparing of the familiarity component of recognition memory in a patient with hippocampal pathology. *Neuropsychologia*, *43*(12), 1810-1823.

Aguirre, G. K., Zarahn, E., & D'esposito, M. (1998). An area within human ventral cortex sensitive to "building" stimuli: evidence and implications. *Neuron*, *21*(2), 373-383.

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

Barense, M. D., Groen, I. I., Lee, A. C., Yeung, L. K., Brady, S. M., Gregori, M., ... & Henson, R. N. (2012). Intact memory for irrelevant information impairs perception in amnesia. *Neuron*, *75*(1), 157-167.

Barr, R. A. et Kaplan LJ (1987). Category representations and their implications for category structure. *Memory and Cognition*, *15*(5), 397-418.

Bartko, S. J., Winters, B. D., Cowell, R. A., Saksida, L. M., & Bussey, T. J. (2007). Perirhinal cortex resolves feature ambiguity in configural object recognition and perceptual oddity tasks. *Learning & Memory*, *14*(12), 821-832.

Behrmann, M., & Avidan, G. (2005). Congenital prosopagnosia: face-blind from birth. *Trends in cognitive sciences*, *9*(4), 180-187.

Bowles, B., Crupi, C., Mirsattari, S. M., Pigott, S. E., Parrent, A. G., Pruessner, J. C., ... & Köhler, S. (2007). Impaired familiarity with preserved recollection after anterior temporal-lobe resection that spares the hippocampus. *Proceedings of the National Academy of Sciences*, *104*(41), 16382-16387.

Brown, M. W., & Xiang, J. Z. (1998). Recognition memory: neuronal substrates of the judgement of prior occurrence. *Progress in neurobiology*, *55*(2), 149-189.

Bruffaerts, R., Dupont, P., Peeters, R., De Deyne, S., Storms, G., & Vandenberghe, R. (2013). Similarity of fMRI activity patterns in left perirhinal cortex reflects semantic similarity between words. *Journal of Neuroscience*, *33*(47), 18597-18607.

Buckley, M. J., & Gaffan, D. (1998). Perirhinal cortex ablation impairs visual object identification. *Journal of Neuroscience*, *18*(6), 2268-2275.

Bussey, T. J., Saksida, L. M., & Murray, E. A. (2002). Perirhinal cortex resolves feature ambiguity in complex visual discriminations. *European Journal of Neuroscience*, *15*(2), 365-374.

Bussey, T. J., Saksida, L. M., & Murray, E. A. (2003). Impairments in visual discrimination after perirhinal cortex lesions: testing 'declarative'vs.'perceptual-mnemonic'views of perirhinal cortex function. *European Journal of Neuroscience*, *17*(3), 649-660.

Capitani, E., Laiacona, M., Mahon, B., & Caramazza, A. (2003). What are the facts of semantic category-specific deficits? A critical review of the clinical evidence. *Cognitive Neuropsychology*, *20*(3-6), 213-261.

Chao, L. L., Haxby, J. V., & Martin, A. (1999). Attribute-based neural substrates in temporal cortex for perceiving and knowing about objects. *Nature neuroscience*, *2*(10), 913-919.

Cipolotti, L., Bird, C., Good, T., Macmanus, D., Rudge, P., & Shallice, T. (2006). Recollection and familiarity in dense hippocampal amnesia: A case study. *Neuropsychologia*, *44*(3), 489-506.

Claessen, M. H., & van der Ham, I. J. (2017). Classification of navigation impairment: A systematic review of neuropsychological case studies. *Neuroscience & Biobehavioral Reviews*, *73*, 81-97.

Clark, R. E., Reinagel, P., Broadbent, N. J., Flister, E. D., & Squire, L. R. (2011). Intact performance on feature-ambiguous discriminations in rats with lesions of the perirhinal cortex. *Neuron*, *70*(1), 132-140.

Clarke, A., & Tyler, L. K. (2014). Object-specific semantic coding in human perirhinal cortex. *Journal of Neuroscience*, *34*(14), 4766-4775.

Cohen, L., Dehaene, S., Naccache, L., Lehéricy, S., Dehaene-Lambertz, G., Hénaff, M. A., & Michel, F. (2000). The visual word form area: spatial and temporal characterization of an initial stage of reading in normal subjects and posterior split-brain patients. *Brain*, *123*(2), 291-307.

Cowell, R. A., Bussey, T. J., & Saksida, L. M. (2006). Why does brain damage impair memory? A connectionist model of object recognition memory in perirhinal cortex. *Journal of Neuroscience*, *26*(47), 12186-12197.

Cowell, R. A., Bussey, T. J., & Saksida, L. M. (2010). Components of recognition memory: dissociable cognitive processes or just differences in representational complexity?. *Hippocampus*, *20*(11), 1245-1262.

Cowell, R. A. (2012). Computational models of perirhinal cortex function. *Hippocampus*, *22*(10), 1952-1964.

Cree, G. S., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*, *132*(2), 163.

Davachi, L., Mitchell, J. P., & Wagner, A. D. (2003). Multiple routes to memory: distinct medial temporal lobe processes build item and source memories. *Proceedings of the National Academy of Sciences*, *100*(4), 2157-2162.

Davachi, L. (2006). Item, context and relational episodic encoding in humans. *Current opinion in neurobiology*, *16*(6), 693-700.

Diana, R. A., Yonelinas, A. P., & Ranganath, C. (2007). Imaging recollection and familiarity in the medial temporal lobe: a three-component model. *Trends in cognitive sciences*, *11*(9), 379-386.

Diana, R. A., Yonelinas, A. P., & Ranganath, C. (2008). High-resolution multi-voxel pattern analysis of category selectivity in the medial temporal lobes. *Hippocampus*, *18*(6), 536-541.

Diana, R. A., Yonelinas, A. P., & Ranganath, C. (2010). Medial temporal lobe activity during source retrieval reflects information type, not memory strength. *Journal of cognitive neuroscience*, *22*(8), 1808-1818.

Downing, P. E., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, *293*(5539), 2470-2473.

Eacott, M. J., & Norman, G. (2004). Integrated memory for object, place, and context in rats: a possible model of episodic-like memory?. *Journal of Neuroscience*, *24*(8), 1948-1953.

Eichenbaum, H., Schoenbaum, G., Young, B., & Bunsey, M. (1996). Functional organization of the hippocampal memory system. *Proceedings of the National Academy of Sciences, 93*(24), 13500-13507.

Eichenbaum, H., Yonelinas, A. P., & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annu. Rev. Neurosci.*, *30*, 123-152.

Epstein, R., Harris, A., Stanley, D., & Kanwisher, N. (1999). The parahippocampal place area: recognition, navigation, or encoding?. *Neuron*, *23*(1), 115-125.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature, 392*(6676), 598-601.

Erez, J., Cusack, R., Kendall, W., & Barense, M. D. (2015). Conjunctive coding of complex object features. *Cerebral Cortex*, *26*(5), 2271-2282.

Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature neuroscience*, *3*(2), 191-197.

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

Graham, K. S., Barense, M. D., & Lee, A. C. (2010). Going beyond LTM in the MTL: a synthesis of neuropsychological and neuroimaging findings on the role of the medial temporal lobe in memory and perception. *Neuropsychologia*, *48*(4), 831-853.

Grill-Spector, K., Knouf, N., & Kanwisher, N. (2004). The fusiform face area subserves face perception, not generic within-category identification. *Nature neuroscience*, *7*(5), 555-562.

Grill-Spector, K., Sayres, R., & Ress, D. (2006). High-resolution imaging reveals highly selective nonface clusters in the fusiform face area. *Nature neuroscience*, *9*(9).

Grüter, T., Grüter, M., & Carbon, C. C. (2008). Neural and genetic foundations of face recognition and prosopagnosia. *Journal of Neuropsychology*, *2*(1), 79-97.

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science*, *293*(5539), 2425-2430.

Huffman, D. J., & Stark, C. E. (2014). Multivariate pattern analysis of the human medial temporal lobe revealed representationally categorical cortex and representationally agnostic hippocampus. *Hippocampus*, *24*(11), 1394-1403.

Jackson, O., & Schacter, D. L. (2004). Encoding activity in anterior medial temporal lobe supports subsequent associative recognition. *Neuroimage*, *21*(1), 456-462.

Johnson-Frey, S. H. (2004). The neural bases of complex tool use in humans. *Trends in cognitive sciences*, *8*(2), 71-78.

Julian, J. B., Ryan, J., & Epstein, R. A. (2016). Coding of Object Size and Object Category in Human Visual Cortex. *Cerebral Cortex*, bhw150.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of neuroscience*, *17*(11), 4302-4311.

Kim, S., Jeneson, A., van der Horst, A. S., Frascino, J. C., Hopkins, R. O., & Squire, L. R. (2011). Memory, visual discrimination performance, and the human hippocampus. *Journal of Neuroscience, 31*(7), 2624-2629.

Kirwan, C. B., & Stark, C. E. (2004). Medial temporal lobe activation during encoding and retrieval of novel face-name pairs. *Hippocampus*, *14*(7), 919-930.

Kivisaari, S. L., Tyler, L. K., Monsch, A. U., & Taylor, K. I. (2012). Medial perirhinal cortex disambiguates confusable objects. *Brain*, *135*(12), 3757-3769.

Konkle, T., & Oliva, A. (2012). A real-world size organization of object responses in occipitotemporal cortex. *Neuron*, *74*(6), 1114-1124.

Konkle, T., & Caramazza, A. (2013). Tripartite organization of the ventral stream by animacy and object size. *Journal of Neuroscience*, *33*(25), 10235-10242.

Kravitz, D. J., Saleem, K. S., Baker, C. I., Ungerleider, L. G., & Mishkin, M. (2013). The ventral visual pathway: an expanded neural framework for the processing of object quality. *Trends in cognitive sciences*, *17*(1), 26-49.

Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis–connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, *2*.

Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., ... & Bandettini, P. A. (2008). Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, *60*(6), 1126-1141.

LaRocque, K. F., Smith, M. E., Carr, V. A., Witthoft, N., Grill-Spector, K., & Wagner, A. D. (2013). Global similarity and pattern separation in the human medial temporal lobe predict subsequent memory. *Journal of Neuroscience, 33*(13), 5466-5474.

Lee, A. C., Buckley, M. J., Gaffan, D., Emery, T., Hodges, J. R., & Graham, K. S. (2006). Differentiating the roles of the hippocampus and perirhinal cortex in processes beyond long-term declarative memory: a double dissociation in dementia. *Journal of Neuroscience*, *26*(19), 5198-5203.

Lee, A. C., Yeung, L. K., & Barense, M. D. (2012). The hippocampus and visual perception. *Frontiers in human neuroscience*, *6*.

Liang, J. C., Wagner, A. D., & Preston, A. R. (2012). Content representation in the human medial temporal lobe. *Cerebral Cortex*, *23*(1), 80-96.

Litman, L., Awipi, T., & Davachi, L. (2009). Category-specificity in the human medial temporal lobe cortex. *Hippocampus*, *19*(3), 308-319.

Maass, A., Berron, D., Libby, L. A., Ranganath, C., & Düzel, E. (2015). Functional subregions of the human entorhinal cortex. *Elife*, *4*, e06426.

Magri, C., Konkle, T., & Caramazza, A. (2016). Visual object responses of the ventral stream reflect both size and motor-relevance. *Journal of Vision*, *16*(12), 505-505.

Mahon, B. Z., & Caramazza, A. (2009). Concepts and categories: A cognitive neuropsychological perspective. *Annual review of psychology*, *60*, 27-51.

Mahon, B. Z., & Caramazza, A. (2011). What drives the organization of object knowledge in the brain?. *Trends in cognitive sciences*, *15*(3), 97-103.

Martin, A., & Chao, L. L. (2001). Semantic memory and the brain: structure and processes. *Current opinion in neurobiology*, *11*(2), 194-201.

Martin, A. (2007). The representation of object concepts in the brain. *Annu. Rev. Psychol.*, *58*, 25-45.

Martin, C. B., McLean, D. A., O'Neil, E. B., & Köhler, S. (2013). Distinct familiarity-based response patterns for faces and buildings in perirhinal and parahippocampal cortex. *Journal of Neuroscience*, *33*(26), 10915-10923.

Martin, C. B., Sullivan, J. A., Wright, J., & Köhler, S. (2018). How landmark suitability shapes recognition memory signals for objects in the medial temporal lobes. *NeuroImage*, *166*, 425-436.

McRae, K., & Cree, G. S. (2002). Factors underlying category-specific semantic deficits. *Category-specificity in brain and mind*, 211-249.

McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior research methods*, *37*(4), 547-559.

McTighe, S. M., Cowell, R. A., Winters, B. D., Bussey, T. J., & Saksida, L. M. (2010). Paradoxical false memory for objects after brain damage. *Science*, *330*(6009), 1408-1410.

Meunier, M., Bachevalier, J., Mishkin, M., & Murray, E. A. (1993). Effects on visual recognition of combined and separate ablations of the entorhinal and perirhinal cortex in rhesus monkeys. *Journal of Neuroscience*, *13*(12), 5418-5432.

Milner, A. D., & Goodale, M. A. (2008). Two visual systems re-viewed. *Neuropsychologia*, *46*(3), 774-785.

Mishkin, M., Ungerleider, L. G., & Macko, K. A. (1983). Object vision and spatial vision: two cortical pathways. *Trends in neurosciences*, *6*, 414-417.

Mishkin, M., & Ungerleider, L. G. (1982). Contribution of striate inputs to the visuospatial functions of parieto-preoccipital cortex in monkeys. *Behavioural brain research*, *6*(1), 57-77.

Mullally, S. L., & Maguire, E. A. (2011). A new role for the parahippocampal cortex in representing space. *Journal of Neuroscience*, *31*(20), 7441-7449.

Mumby, D. G., & Pinel, J. P. (1994). Rhinal cortex lesions and object recognition in rats. *Behavioral neuroscience*, *108*, 11-11.

Mumby, D. G. (2001). Perspectives on object-recognition memory following hippocampal damage: lessons from studies in rats. *Behavioural brain research*, *127*(1), 159-181.

Mumby, D. G., Gaskin, S., Glenn, M. J., Schramek, T. E., & Lehmann, H. (2002). Hippocampal damage and exploratory preferences in rats: memory for objects, places, and contexts. *Learning & Memory*, *9*(2), 49-57.

Murray, E. A., & Bussey, T. J. (1999). Perceptual–mnemonic functions of the perirhinal cortex. *Trends in cognitive sciences*, *3*(4), 142-151.

Murray, E. A., & Richmond, B. J. (2001). Role of perirhinal cortex in object perception, memory, and associations. *Current opinion in neurobiology*, *11*(2), 188-193.

Murray, E. A., Bussey, T. J., & Saksida, L. M. (2007). Visual perception and memory: a new view of medial temporal lobe function in primates and rodents. *Annu. Rev. Neurosci.*, *30*, 99-122.

Nemanic, S., Alvarado, M. C., & Bachevalier, J. (2004). The hippocampal/parahippocampal regions and recognition memory: insights from visual paired comparison versus object-delayed nonmatching in monkeys. *Journal of Neuroscience*, *24*(8), 2013-2026.

Proklova, D., Kaiser, D., & Peelen, M. V. (2016). Disentangling representations of object shape and object category in human visual cortex: The animate–inanimate distinction. *Journal of cognitive neuroscience*.

Ranganath, C., Yonelinas, A. P., Cohen, M. X., Dy, C. J., Tom, S. M., & D'Esposito, M. (2004). Dissociable correlates of recollection and familiarity within the medial temporal lobes. *Neuropsychologia*, *42*(1), 2-13.

Ranganath, C., Cohen, M. X., Dam, C., & D'Esposito, M. (2004). Inferior temporal, prefrontal, and hippocampal contributions to visual working memory maintenance and associative memory retrieval. *Journal of Neuroscience*, *24*(16), 3917-3925.

Saksida, L. M., & Bussey, T. J. (2010). The representational–hierarchical view of amnesia: Translation from animal to human. *Neuropsychologia*, *48*(8), 2370-2384.

Sperling, R., Chua, E., Cocchiarella, A., Rand-Giovannetti, E., Poldrack, R., Schacter, D. L., & Albert, M. (2003). Putting names to faces:: Successful encoding of associative memories activates the anterior hippocampal formation. *Neuroimage*, *20*(2), 1400-1410.

Squire, L. R., & Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science*, *253*(5026), 1380.

Squire, L. R., Stark, C. E., & Clark, R. E. (2004). The medial temporal lobe. *Annu. Rev. Neurosci.*, *27*, 279-306.

Staresina, B. P., & Davachi, L. (2006). Differential encoding mechanisms for subsequent associative recognition and free recall. *Journal of Neuroscience*, *26*(36), 9162-9172.

Suzuki, W. A., & Amaral, D. G. (2004). Functional neuroanatomy of the medial temporal lobe memory system. *Cortex*, *40*(1), 220-222.

Suzuki, W. A. (2009). Perception and the medial temporal lobe: evaluating the current evidence. *Neuron*, *61*(5), 657-666.

Suzuki, W. A., & Baxter, M. G. (2009). Memory, perception, and the medial temporal lobe: a synthesis of opinions. *Neuron*, *61*(5), 678-679.

Takahashi, N., & Kawamura, M. (2002). Pure topographical disorientation—the anatomical basis of landmark agnosia. *Cortex*, *38*(5), 717-725.

Taylor, K. I., Devereux, B. J., Acres, K., Randall, B., & Tyler, L. K. (2012). Contrasting effects of feature-based statistics on the categorisation and basic-level identification of visual objects. *Cognition*, *122*(3), 363-374.

Thomas, C., Avidan, G., Humphreys, K., Jung, K. J., Gao, F., & Behrmann, M. (2009). Reduced structural connectivity in ventral visual cortex in congenital prosopagnosia. *Nature neuroscience*, *12*(1), 29-31.

Troiani, V., Stigliani, A., Smith, M. E., & Epstein, R. A. (2012). Multiple object properties drive scene-selective regions. *Cerebral Cortex*, *24*(4), 883-897.

Tyler, L. K., Chiu, S., Zhuang, J., Randall, B., Devereux, B. J., Wright, P., ... & Taylor, K. I. (2013). Objects and categories: feature statistics and object processing in the ventral stream. *Journal of Cognitive Neuroscience*, *25*(10), 1723-1735.

Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, *107*(3), 829-853.

Yassa, M. A., & Stark, C. E. (2008). Multiple signals of recognition memory in the medial temporal lobe. *Hippocampus*, *18*(9), 945-954.

Yonelinas, A. P. (2002). The nature of recollection and familiarity: A review of 30 years of research. *Journal of memory and language*, *46*(3), 441-517.

Chapter 2

2 Animacy and real-world size shape object representations in the medial temporal lobes

2.1 Introduction

The ability to identify the objects we encounter in our daily lives, and know which ones we have seen before, is a crucial aspect of human behavior. What type of object a ‘thing’ is, and whether it is familiar or novel, can drastically change how we might interact with it, including, for example, whether to approach or avoid it. Despite the importance of object recognition, and the relative fluidity with which most humans perform it, a detailed understanding of the neural functional architecture that supports this ability is still elusive. One promising approach to understanding the neural architecture of object perception and memory is to explore how object representations are organized. In particular, it is possible to examine similarities between patterns of brain activity that different types of objects evoke, and to map this neural representational space to relevant dimensions in perception and behavior.

It is known that some correspondence exists between how objects are represented in the brain and how we behaviorally categorize them. Crucial evidence comes from neuropsychological research in patients with focal brain lesions who have category specific behavioral deficits in object processing (Warrington & Shallice, 1984; Hart et al., 1985; for a review see Caramazza & Mahon, 2003). Interestingly, patients have been reported who have impairments for a broad domain of object types, such as either living or nonliving objects (for a review see Capitani et al., 2003), while other patients have shown deficits for specific categories that fall within these domains, such as faces or tools (Damasio et al., 1982; De Renzi et al., 1994; for a review see Johnson-Frey, 2004). Loss of category specific knowledge, and the more common loss of domain knowledge beg questions of how objects from different categories and domains are represented in the brain.

Some initial answers to these questions have come from functional magnetic resonance imaging (fMRI) investigations of object processing in the ventral visual stream (VVS). Numerous fMRI studies have revealed regions within the VVS that preferentially respond to particular stimulus categories with high ecological relevance, including faces, scenes, bodies, and words (see Op de Beeck et al., 2008, for review). Specifically, in extrastriate cortex, distinct circumscribed regions have been reported that prefer one of these categories over other categories, such as the fusiform face area or the parahippocampal place area (Kanwisher et al., 1997; Epstein & Kanwisher, 1998; Downing et al., 2001; for a review see Kanwisher & Dilks, 2013). Interestingly, these functionally circumscribed regions are systematically organized within broader preference zones. Medial aspects of occipito-temporal cortex typically show a preference for inanimate objects, whereas lateral aspects show a preference for animate objects (Martin, 2007; Grill-Spector & Weiner, 2014; Sha et al., 2015). In addition to the animacy dimension, a number of fMRI studies have revealed large-scale organization of the VVS has also been linked to real-world size (Konkle & Caramazza, 2013; Konkle, et al., 2012; Mitchell & Cusack, 2016). It has been found that there is a preference zone for large inanimate objects in medial occipito-temporal cortex and for small inanimate objects in more dorsolateral aspects, but no corresponding size-based distinction has been found for animate objects in lateral occipito-temporal cortex. This pattern of preferences has sometimes been referred to as a tripartite organizing schema (Konkle & Oliva, 2012; Konkle & Caramazza, 2013).

These findings, that some regions show higher levels of activity on average for one category over another (e.g., FFA), or for one domain over another (e.g., medial VVS for inanimate objects) have primarily been mapped in fMRI studies by contrasting averaged BOLD activity for one category or domain of stimuli with another. Notably, however, revealing differential activity using this approach does not necessarily imply that a given region *only* represents stimuli from categories or domains to which it is maximally responsive. In fact this has been shown not to be the case within the FFA, where it is possible to detect differentiation between non-face categories of objects using multivariate pattern analysis (MVPA) (Haxby et al., 2001). In this approach, activity is not averaged across voxels, but the similarity between patterns of activity evoked by different stimuli within a given region are compared. If stimuli within a category evoke more similar patterns of activity than stimuli from different categories, the brain region is considered to contain representations of that category. Inasmuch as the pattern of activity across voxels can be labeled a neural representation of an object, one can think of the comparisons between categories as now existing in “representational space”. Interestingly, Kriegeskorte et al., (2008) applied this approach across the entire ventral temporal cortex for a wide variety of objects, and found a highly consistent category- and domain-based organization, with evidence for a distinction between animate and inanimate objects, as well as varying degrees of similarity between categories within these domains (see also Proklova, 2016). A recent fMRI study with a similar focus on representational similarities has shown that real-world size is also an organizing dimension of objects across a large swath of temporo-parieto-occipital cortex, as well as within a number of subregions across the VVS (Julian et al., 2016).

Category and domain preferences have been most thoroughly characterized in the posterior and lateral aspects of the VVS. At present, evidence that speaks to the organization of object representations in medial temporal structures, specifically perirhinal cortex (PrC) and parahippocampal cortex (PhC), is more limited. To be more precise, the more posterior aspect of the PhC has been well characterized, given that it comprises a significant proportion of the parahippocampal place area, a functionally-defined region that preferentially responds to scenes and large objects, in particular those with navigational relevance (Aguirre et al., 1998; Epstein and Kanwisher, 1998; Epstein & Vass, 2014; Troiani et al., 2012; Konkle & Oliva, 2012).

However, it is less clear whether this characterization holds for PhC as a whole, and, precisely how object representations in PrC differ from those in PhC in terms of organization. The lack of evidence is surprising given that PrC has been proposed to be the apex of the VVS (Murray and Bussey, 1999; Bussey et al., 2007). In particular, it has been argued that PrC contains high-level visual object representations, composed of conjunctions of lower-level visual feature information that converges from downstream visual areas. According to this representational-hierarchy view, these high-level visual object representations tend to be particularly relevant for the recognition that an object has been seen before, because they correspond to the types of things in the world for which humans and other animals form distinct declarative memories (e.g., a specific person's face). However, they are also thought to be crucial for visual perception to the extent that the task at hand requires fine-grained discrimination between objects with many overlapping features (e.g., distinguishing a horse from a donkey). The characterization of object representations in PhC, primarily in the parahippocampal place area (PPA), by contrast, has focused less on the role of feature conjunctions and instead addressed the sensitivity to specific types of visual features, such as texture, spatial frequency, and real-world size (Cant & Goodale, 2007; Park & Park, 2017; Rajimehr, 2011; Julian et al., 2016; Konkle and Oliva, 2012) as well as sensitivity to functional aspects, such as whether objects are space-defining, or have navigational relevance (Janzen & van Turennout, 2004; Mullaly & Maguire, 2011; Troiani et al.,2012; Martin et al., under review). A direct comparison of representational space for PrC and PhC across many object categories or larger domains is missing at present. Specifically, it is not known whether the major dimensions that have been shown to shape representations in the posterior VVS, i.e., animacy and real-world size, also shape organization of object categories in PrC and PhC.

Anatomically, evidence from studies of structural connectivity in non-human primates as well as functional connectivity studies in humans suggest that both PhC and PrC have strong connectivity with downstream areas in the VVS and other posterior cortical regions (see Ranganath & Ritchey, 2012, for a review). Interestingly, these connections are at least in part distinct. Both tract tracing data from monkeys, and functional connectivity data in humans show the following differential connectivity between the two regions. PhC is highly connected to earlier occipital and temporal areas as well as densely connected to the retrosplenial cortex (RsC), whereas PrC is more densely connected to higher order regions within the ventral temporal cortex, such as the anterior fusiform gyrus. PhC and PrC also have differential connectivity patterns with structures outside of the VVS. Specifically, PhC is connected to the medial parietal cortex, precuneus, ventrolateral parietal cortex and medial prefrontal cortex (default mode areas), while PrC is more densely connected with the amygdala, orbitofrontal cortex, and anterior ventrolateral temporal cortex (Kahn et al., 2008; Libby et al., 2012). One possibility is that the differential structural and functional connectivity of PrC and PhC leads to differences in the organization of objects representations in both structures.

Research with direct comparisons of visual stimulus responses in PrC and PhC has shown robust differences for processing of faces, objects, and scenes across both structures. At the univariate level, PhC shows a scene preference, while PrC, in particular anterior portions, show a face preference (Liang et al., 2012; Litman et al., 2010; O'Neil et al., 2013; see Collins and Olson, 2014, for review). In MVPA based studies it has been shown that object, scene, and face information can be distinguished at the category level in both PhC and PrC. In general scene decoding is much higher in PhC, and face responses can be better decoded from PrC (LaRocque et al., 2013; Huffman & Stark, 2014, Liang et al., 2012), although Diana et al. 2008, did not find above chance decoding of objects or faces in PrC. While there is clearly a distinction between face and scene representations across both structures, it is less clear whether there are also differences in representations of different types of objects categories. This is in large part due to the fact that most studies probed mixed groups of objects without any systematic attempt to probe category based distinctions. In recent work from our lab, Martin et al. (2013; 2016) explored this issue in the context of recognition memory judgments, using chairs, faces, and buildings as categorized stimuli. We reported that it was possible to decode the perceived familiarity of faces from activity patterns in PrC, the familiarity of buildings from patterns in PhC, and familiarity for chairs from patterns in both structures. While these findings go beyond showing a distinctions between scenes and faces in the MTL, they do not allow for a broader characterization of representational space across a wider variety of object categories.

Our primary interest in the present study was in a comparison of object representations across the PrC, PhC, and in the hippocampus (HpC). While previous work shows some object category specificity in PrC and PhC, the hippocampus has been seen as more "agnostic", or insensitive to visual stimulus category (Huffman & Stark, 2014; LaRocque et al., 2013; Diana et al., 2009). It has been posited that this is because the HpC binds together object and spatial information received from the PrC and PhC (Eichenbaum et al., 2012; Ranganath & Ritchey, 2012). More specifically, if the HpC represents complex conjunctions of many different kinds of objects and the spatial backdrop of those objects, it may be difficult to reveal any category specificity (e.g., in a complex scene there may be objects from many different categories). Interestingly, one study reported above-chance decoding of scene information from posterior HpC (Liang et al., 2012). This selective above chance decoding of scene stimuli converges with a body of research suggesting the HpC is differentially involved in scene processing (Hodgett's et al., 2016; Lee et al., 2005; Barense et al., 2015; Zeidman et al., 2015,; for a review see Murray et al., 2017). At the univariate level, the HpC often shows more activity for scenes as compared to other stimulus categories such as objects or faces, which has led to the suggestion it be considered a part of the core scene-network (Hodgetts et al., 2016). Furthermore, individuals with hippocampal damage have been reported to show impairments in vividly recalling scenes, maintaining scenes in working memory, or constructing scenes in their imagination (Hassabis et al., 2007; Mullally et al., 2012; Addis et al., 2007). Additionally, distinct hippocampal activation patterns have been reported in response to these different aspects of scene processing in fMRI research (Zeidman et al., 2015). This suggests that the HpC may not be entirely agnostic to the nature of stimulus categories encountered. As such, it is possible that it may also be sensitive to stimulus domain.

In the present fMRI study, we directly aimed to address whether and how animacy and real-world size affect the organization of object categories in the MTL. We tested the hypothesis that object-evoked responses in perirhinal and parahippocampal cortex show evidence for domain-level organization along both dimensions. To this end, we scanned participants while they performed a continuous recognition memory task on objects from 12 different categories. We chose a continuous recognition memory task because it required participants to make memory decisions (i.e., "old" or "new") for specific exemplars from these categories, thus maximizing the need to disambiguate objects with substantial feature overlap. To address our questions of interest, we employed representational similarity analyses (RSA). With these analyses we first asked whether PrC, PhC and the HpC represent distinct categories of objects. We then explored whether the categories were organized along an animate/inanimate divide, and whether or not inanimate objects were organized by their real-world size.

.

2.2 Materials and methods

2.2.1 Participants

Fifteen individuals participated in the study (20-32 years of age, mean age = 27.5 years; 8 females). All participants were right-handed with normal or corrected-to-normal vision, and no history of psychiatric or neurological disorders. Data from two participants were excluded due to technical difficulties. Participants received financial compensation for their participation, and provided informed consent according to procedures approved by the University of Western Ontario Health Sciences Research Ethics Board

2.2.2 Stimuli

Stimuli were color images depicting exemplars from 12 different object categories, including 4 categories of animate objects (faces, bodies, monkeys, insects), 4 categories of large inanimate objects (buildings, vehicles, trees, furniture), and 4 categories of small inanimate objects (flowers, fruits, musical instruments, and tools). Size and animacy classification was based on prior research (Konkle et al., 2013) and confirmed through ratings in pilot work in a separate group of participants for all stimuli employed here. Twenty-eight objects were chosen from each category, for a total of 336 experimental stimuli. In addition, 3 filler items were presented in each run, one of which was repeated early on in the run to ensure that participants would immediately be prepared for repetitions. The second and third filler items were presented towards the end of the run to increase the proportion of novel stimuli at that stage. Filler items were chosen from categories other than (and unrelated to) those employed on experimental trials. Images of objects were obtained from the Konkle lab database (http://konklab.fas.harvard.edu/#) and through an additional Google image search. Each image was presented in isolation on a uniform grey background. The size of each image was bound at a maximum of 500 x 500 pixels, with at least one dimension exactly corresponding to these limits. Across categories, there were no significant differences in the area covered by objects in the images, their aspect-ratio, or their mean luminance (all $p > 0.05$).

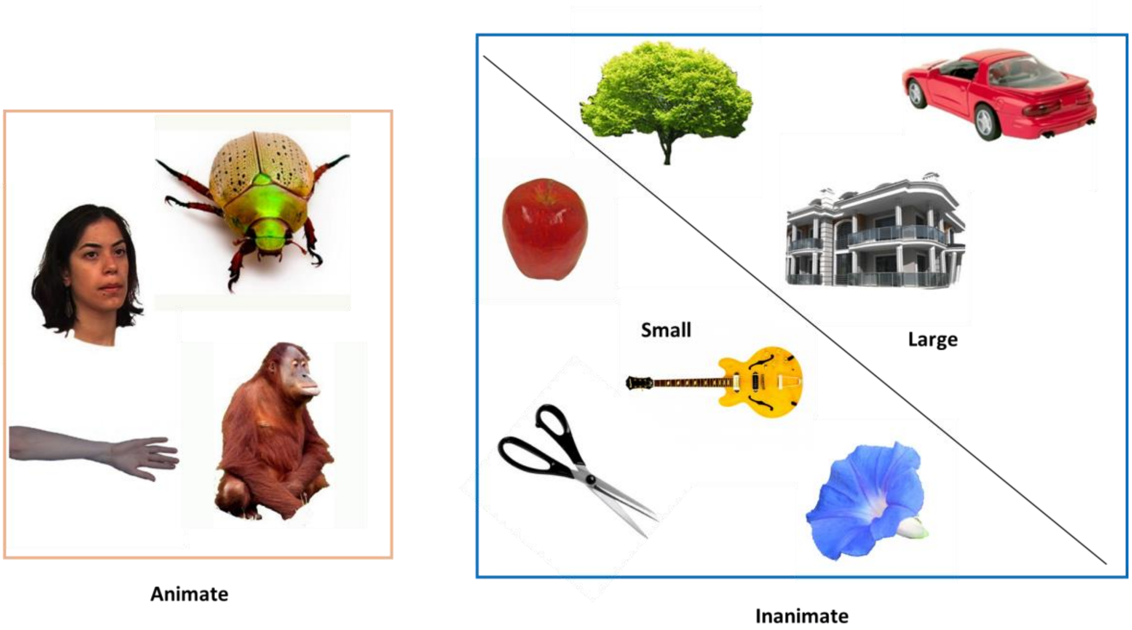

Figure 2.1: Stimuli from 12 different object categories, separated by animacy and real-world size for inanimate objects. Categories were grouped into animate: faces, bodies, monkeys, insects, inanimate small: flowers, fruits, tools, musical instruments, and inanimate large: buildings, trees, vehicles, furniture.

2.2.3 Experimental procedure

During fMRI scanning, participants performed a continuous recognition memory task that required recognition of repeated presentation of specific exemplars (Figure 2.2). Exemplars were presented twice, with repetitions always occurring in the same run. Images were presented for 1200 ms, and participants were asked to indicate whether the image was “novel” (1st presentation), or “old” (2nd presentation) with button presses using their middle or index finger. To encourage rapid responding and mark the time window for responding, a red border surrounding the image appeared 600 ms after stimulus onset and stayed on screen until stimulus offset. Participants were instructed to respond as soon as the red border appeared. Mapping of responses to buttons was counterbalanced across participants. Each stimulus presentation was followed by a jittered ITI (2000-6000 ms) during which participants viewed a fixation cross centered on a grey background. Jitter was distributed such that the average delay between first and second presentation of items was matched across categories (average time = 84.1 s, range = 19.0-316.0 s). In addition, the average number of images between repetitions was matched across categories (average number of intervening images =17, range =16-18*)*. Each run consisted of 4 objects from each of the 12 categories, resulting in a total of 8 image presentations per category, or 96 experimental trials per run. In addition, each run contained 3 filler trials. Across runs, presentations of objects from each category were preceded and followed by an object from each of the other categories with roughly equal frequency (8-11 times*)*. Participants completed seven runs. Three different run orders were created for the purpose of counterbalancing across participants. Prior to scanning, each participant completed a 5-minute practice task with images from categories that were unrelated to those used during scanning in order to be familiarized with task requirements and response deadline.

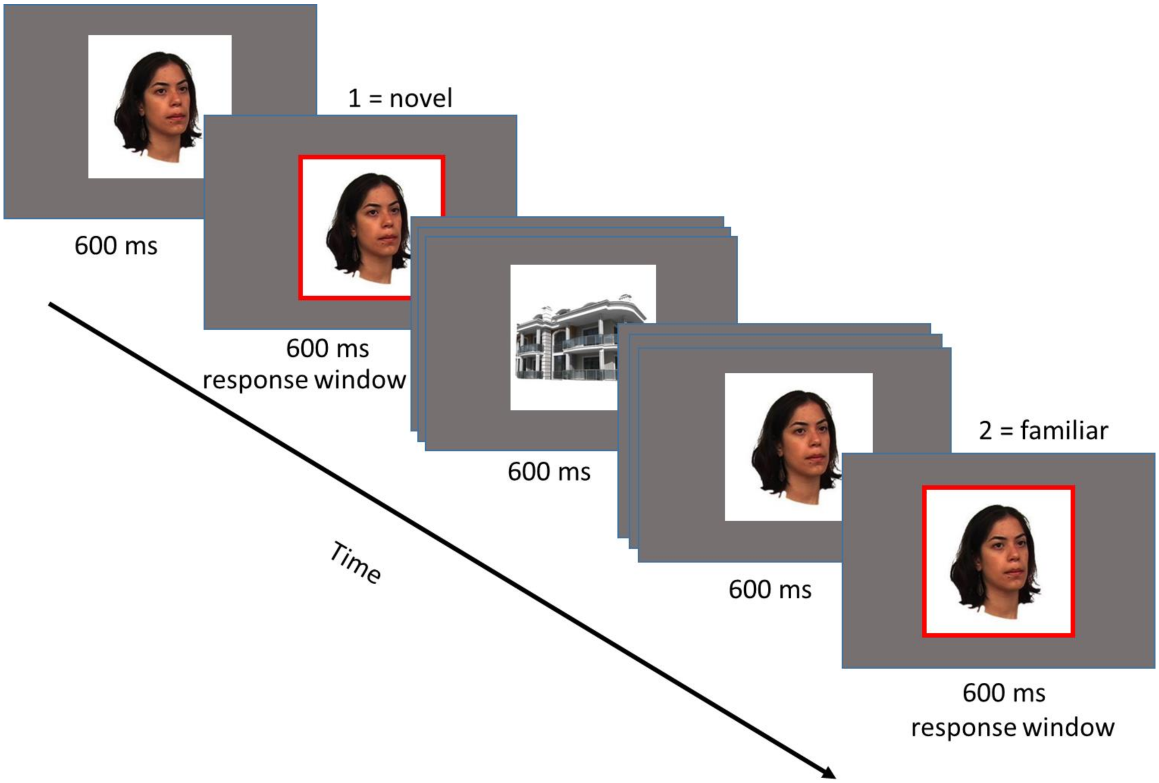

Figure 2.2: Continuous recognition memory task. An image depicting an object from one of the twelve categories was presented on screen for 1200ms total. After 600ms, a red border popped up around the image, and participants had the remaining 600ms to respond "novel" indicating it was the first time they had seen that image, or "familiar" indicating that it was the second time they had seen that image. Each run contained different stimuli, so repetitions were on a within-run basis. There were a number of intervening stimuli between repetitions (on average 20 intervening stimuli).

2.2.4 Image acquisition

MRI data was acquired on a Siemens TIM Trio 3-Tesla scanner. Functional MRI volumes were collected using a highly accelerated gradient-echo EPI sequence (Center for Magnetic Resonance Research, University of Minnesota) with a multiband acceleration factor 4 and GRAPPA in-plane acceleration of 2. The following parameters were used: TR=650 ms, TE=30 ms, slice thickness = 2 mm, FOV = 192 mm X 192 mm, flip angle = 54 degree . Each functional volume included 40 slices collected in an interleaved manner. To optimize MR signal in the anterior temporal lobes, a transverse orientation was chosen for acquisition, which allowed for inclusion of the entire temporal and occipital lobes, with partial coverage of frontal and parietal cortices, in all participants. T1-weighted anatomical images were obtained using an ADNI MPRAGE sequence (192 slices, TR = 2300 ms, TE =2.98 ms, 1 mm isotropic voxels, FOV = 240 X 256 mm, flip angle = 9degrees).

2.2.5 Neuroimaging analyses

2.2.5.1 Pre-processing and modeling

fMRI data were analyzed using SPM8 (Welcome Institute of Cognitive Neurology; http://www.fil.ion.ucl.ac.uk/spm/software/spm8/), employing an analysis pipeline as implemented in the automatic analysis system (aa) (www.github.com/rhodricusack/automaticanalysis), (Cusack et al., 2015). Functional data were motion corrected and high pass-filtered to remove low frequency noise (drift); slice-time correction was not implemented due to the use of a multiband sequence. Four dummy scans at the start of each session were discarded to allow for T1 relaxation. For each participant, the mean functional image was then co-registered with the participant-specific anatomical image. Following co-registration two separate pipelines were initiated: a univariate pipeline and an MVPA pipeline. For the univariate pipeline, co-registered data was normalized to MNI space and smoothed using a three-dimensional Gaussian kernel with a full-width at half maximum of 8 mm. For the MVPA pipeline, co-registered images were kept in native space for each participant, and no spatial smoothing was applied in order to preserve high-spatial resolution. Functional data were convolved using a canonical hemodynamic response function. Categories were modeled, regardless of whether a trial was a 1st or 2nd presentation (12 regressors per run) using a general linear model. Regressors were constructed from boxcars with a durations of each stimulus

(1200ms), and were convolved with SPM's canonical hemodynamic response function. Beta estimates for each category were derived based on 4 exemplars and their repetition in each run. Regressors of no interest included 6 motion regressors. Beta estimates derived from these models were used as input for the univariate and multivariate analyses. Medial temporal lobe ROIs were demarcated manually for each participant on the high-resolution structural images in native space, using the anatomical protocols published by Pruessner et al. (2000; 2002) with adjustments to the posterior border of PhC as specified by Franko et al. (2014).

2.2.5.2 Univariate analyses

Univariate analyses were conducted for feature selection of voxels to be included in the multivariate analyses. Towards this end, we contrasted all experimental trials against baseline (gray screen with a fixation cross), which resulted in robust activation throughout occipital and temporal cortex (including MTL) in each participant. We then contrasted each anatomically defined region of interest in the MTL, and selected the 20% of voxels with the highest beta values in this contrast (i.e., stimuli vs baseline). These voxels for all remaining multivariate analyses (see Kriegeskorte et al., 2008, 2008b for rationale).

2.2.5.3 Representational similarity analysis

Multivariate analyses were computed on a between-run basis to ensure the different comparisons did not vary in temporal proximity (Linke et al., 2011). To explore the representational space in each ROI, for each subject, we first extracted beta values for each category and computed the Pearson's correlation for each category compared to each other category. Prior to computing the correlations, the grand mean (i.e., the cocktail mean) for each run was subtracted across all voxels for that run (Walther et al. 2015). This resulted in a 12 x 12 representational similarity matrix (RSM) for each participant, for each ROI, with within category similarity values (across runs) on the diagonal, and between category information (across runs) on the off diagonal. RSMs were then averaged across participants, resulting in a final group similarity matrix for each ROI. For visualization purposes, group RSMs were ordered in the following way: animate objects, small inanimate objects, and large inanimate objects. Note that RSM's are not symmetrical in the visualization, this is because the upper half of each matrix shows the mean from a subset of across run correlations (i.e., cell 1, 2 is condition in the even runs correlated with condition 2 in the odd runs, whereas cell 2, 1, is condition 1 in the odd runs correlated with condition 2 in the even runs) (Figure 2.4). To test whether the representational space was modulated by category, animacy, and size within inanimate objects, we created linear models (predefined contrasts) specifying which RSM correlation values were to be subjected to a t-test that tested models (see Figure 2.5). These analyses were performed on data in single-subject RSMs, with the group statistics calculated from the average results.

We first asked whether there was evidence of category-level organization in each ROI. To test for this, we defined a contrast of category representation (see Fig. 2.5), in other words, a linear model where all within category (diagonal) patterns were more highly correlated than between category (off diagonal) patterns. In the initial analysis, we tested an omnibus contrast (i.e., model) that probed for the presence of any category-specific information in each ROI. We then tested for information relating to each of the 12 categories individually. Specifically, we tested whether the patterns of activation across voxels were more similar within each category compared to the 11 other categories, using subject as a random effect.

In our second set of analyses, we asked whether or not the animate v. inanimate object distinction that has been found to shape the organization of object representations in more posterior aspects of the VVS was also an organizing dimension in the MTL. This analysis was identical to the previously described analyses, except that for the purpose of evaluating differences in correlations (i.e., within vs between) we focused on the domains of animate as compared to inanimate objects rather than individual categories (see Figure

2.5). Importantly, in these analyses we removed the diagonal from our model in order to discard the influence of within category similarities.

In our third and final set of analyses, we asked whether real world size is an organizing dimension within the domain of inanimate objects in MTL, again as has been reported for object representations in more posterior aspects of the VVS. Here, we divided inanimate objects into groups of small or large objects, large objects included trees, furniture, vehicles and buildings, and small objects including fruit, flowers, musical instruments, and tools. The analysis was identical to the previous one except that within versus between similarities were computed across all categories of large or small inanimate objects (see Figure 2.5). As in the analyses on animacy described above, we did not include the diagonal in testing of this model.

2.3 Results

2.3.1 Behavioural

Recognition-memory accuracy, indexed using the discriminability index *d*', and reaction times are shown in Table 2.1 for all categories. Critically, memory discrimination as measured with *d'* was matched across dimensions of interest. Specifically, we found no differences in performance between animate and inanimate objects (Mean *d'* inanimate = 1.76, *SD*=0.78, Mean *d'* animate =1.94, SD=0.71, $t(12) = -1.30$, $p = .2$ (Figure 2.3)). There were also no differences based on real-world size, i.e., between large inanimate and small inanimate objects (Mean *d'* small inanimate = 1.96, *SD*=0.77, Mean d' large inanimate = 2.00, SD=0.80, $t\ (12) = -.452$, $p = .7$ (Figure 2.3)). We did find differences in RTs between animate and inanimate categories (Mean RT animate = 1.00, *SD* = 0.036, Mean RT inanimate = 1.012 *SD* = 0.036 $t(12) = -2.709$, $p < 0.02$), as well as large inanimate and small inanimate objects (Mean RT large inanimate = 1.007s, *SD* = 0.40, Mean RT small inanimate = 0.993, *SD* = 0.033, $t(12) = 3.413$, $p < 0.005$).

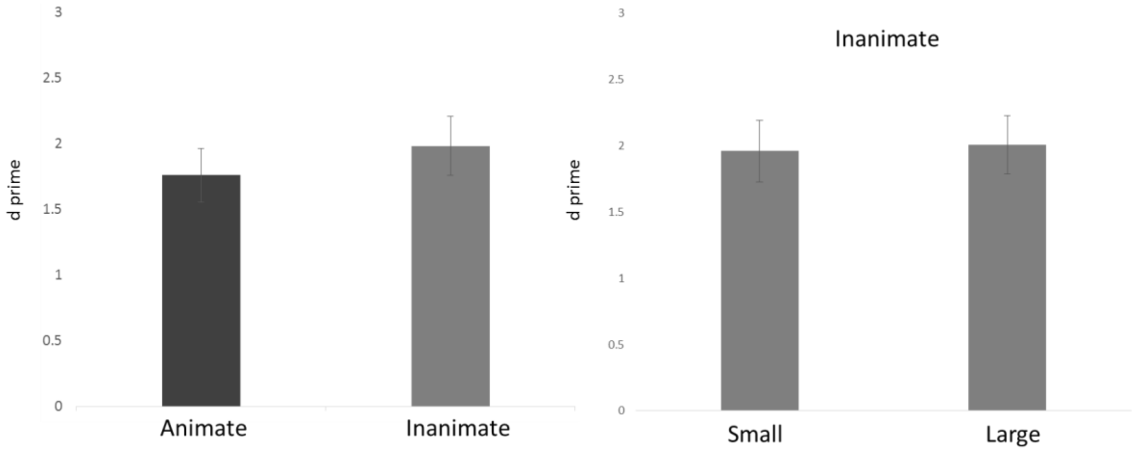

Figure 2.3: Behavioral performance: continuous recognition memory task. Average recognition memory performance on dimensions of interest across all subjects, as measured by *d'*. There were no significant differences in recognition memory between animate and inanimate objects (p = .2), or between large inanimate and small inanimate objects (p = .7).

Recognition memory performance by category

category	**d’**	**RT cr’s**	**RT hits**	**Sig differences overall RT**	**Sig differences dprime**
face	1.58, sd=0.35	1.037, sd=0.041	.975 sd=0.057		
body	1.60, sd=0.61	1.059, sd=0.042	.979, sd=0.045	fruit, tool	
monkey	1.69, sd=0.66	1.056, sd=0.047	.980, sd=0.034	fruit, tool	
insect	2.18, sd=0.97	1.041 sd=0.038	.972, sd=0.046		
flower	1.93, sd=0.86	1.017, sd=0.027	.991, sd=0.070		
fruit	2.33, sd=0.95	1.023, sd=0.045	.951, sd=0.038	tree, building	tree
musical instrument	1.88, sd=0.59	1.012, sd=0.034	.968, sd=0.046		
tool	1.90, sd=0.58	1.014, sd=0.048	.965, sd=0.047	building	
tree	1.62, sd=0.71	1.035, sd=0.042	.994, sd=0.042		
vehicle	2.51, sd=0.86	1.007, sd=0.047	.987, sd=0.063		body, face, fruit, furniture, tree, monkey
furniture	1.71, sd=0.64	1.029, sd=0.042	.974, sd=0.040		
building	1.99, sd=0.72	1.036, sd=0.048	.995, sd=0.049		

Table 2.1: Behavioral performance on the continuous recognition memory task by object category. Significant differences are listed for reaction times (shown here in seconds) pooled (correct rejections and hits) for each object category, as well as for category differences in overall performance as measured by *d’*. cr’s indicates correct rejections.

2.3.2 fMRI

2.3.2.1 Category

We first tested a model that probed for the presence of category-specific information by comparing within versus between category similarity across all categories combined, we employed Bonferroni correction for the number of ROIs (3) (Figure 2.5). We found that all MTL regions showed sensitivity to category membership (PhC: $t(12) = 6.41$, $p = .00006$; PrC: $t(12) = 5.01$, $p = .0006$; HpC: $t(12) = 3.67$, $p = .009$ (Figure 2.5). Next we examined sensitivity to information about each category individually, asking for each category whether the within pattern similarity for that category (across runs) was more similar than the between pattern similarity (for that category compared to all other tested categories across runs). To adjust for the larger number of corresponding comparisons, we employed Bonferroni correction in these analyses. In PhC, we found significant effects for buildings ($t(12) = 5.62$, $p = .001$), furniture ($t(12) = 3.85$, $p = .02$), vehicles ($t(12) = 4.15$, $p = .01$), and faces ($t(12) = 4.23$, $p = .01$). In PrC we found category related effects for monkeys ($t(12) = 4.28$, $p < .01$), and a trend towards significance for faces ($t(12) = 3.17$, $p < 0.08$, uncorrected p = 0.007). In the HpC, we only found one category that showed a trend towards significance, namely buildings ($t(12) = 3.37$, $p < .06$, uncorrected p = 0.005).

2.3.2.2 Animacy

In our next set of analyses we turned to domain-level organization of object representations based on groupings of multiple categories. Specifically, we asked whether MTL regions hold information shared between categories at the domain level of animacy. To address this question, we probed whether representations for objects within a domain (animate or inanimate, respectively) share more similarity with each other than they do with representations from the other domain. In order to remove any impact of category-level effects (as described in the previous section), we removed the diagonal in this model (Figure 2.5). We found that the representational structure in both PhC and PrC reflected the animacy divide *(*PhC: *t*(12) = 3.73, *p* = .002; PrC: *t*(12)= 3.02, *p* = .02). By contrast, we found no evidence for organization of object representations by animacy in the HpC (*t*(12) = 2.04, p = 0.18) (Figure 2.5).

2.3.2.3 Real-world size

In a further set of analyses, we examined domain-level organization related to the size of inanimate objects. To address this question, we probed whether representations for objects within the domain of small or large inanimate objects, respectively, share more similarity with each other than they do with representations from the other domain. Again, we removed the diagonal in this model in order to remove any impact of category-level effects (Figure 2.5). We found evidence for size related organization in both the PhC and HpC (PhC: *t*(12) = 4.14, *p* = .003; HpC: t(12) = 4.07, p = .003). By contrast we found no such evidence in PrC (t(12) =2.67, p = 0.06) (Figure 2.5).

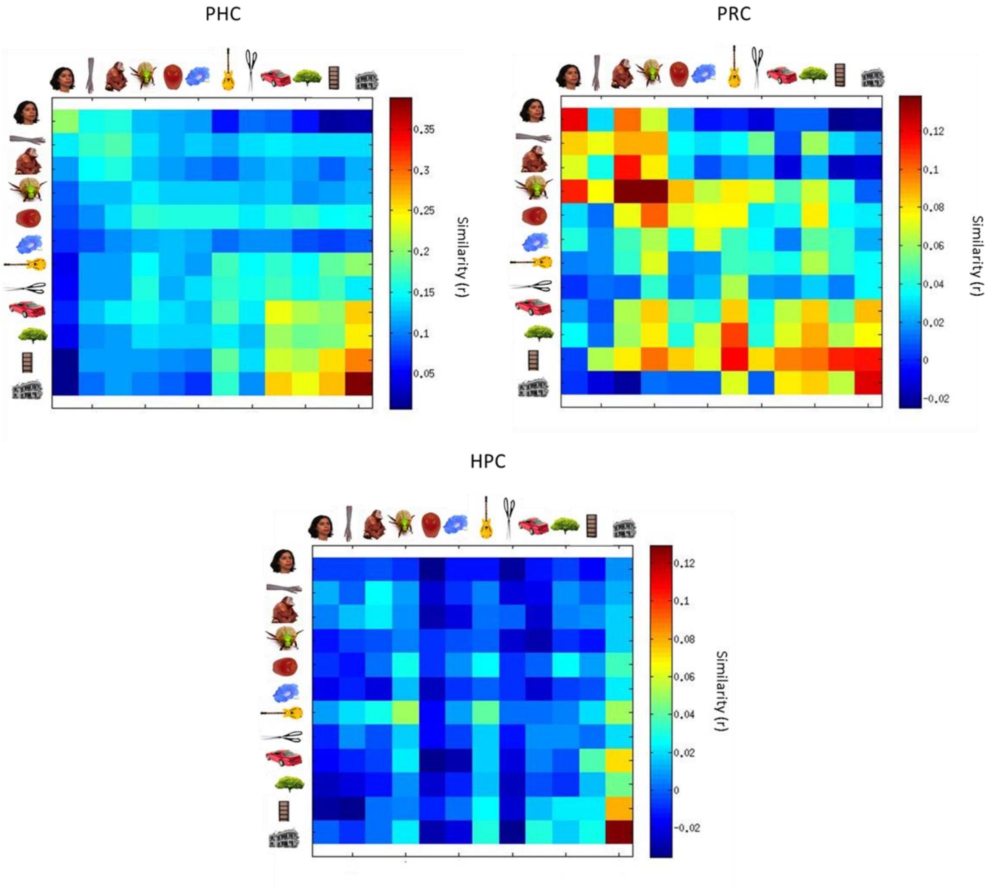

Figure 2.4: Representational space for object-evoked responses in the medial temporal lobe. Representational similarity matrices for the three MTL structures. Matrices show Pearson's correlations between patterns of activity evoked by each object category compared to each other object category. Note that the diagonal shows within-category correlations across runs (each run had different exemplars from the given category).

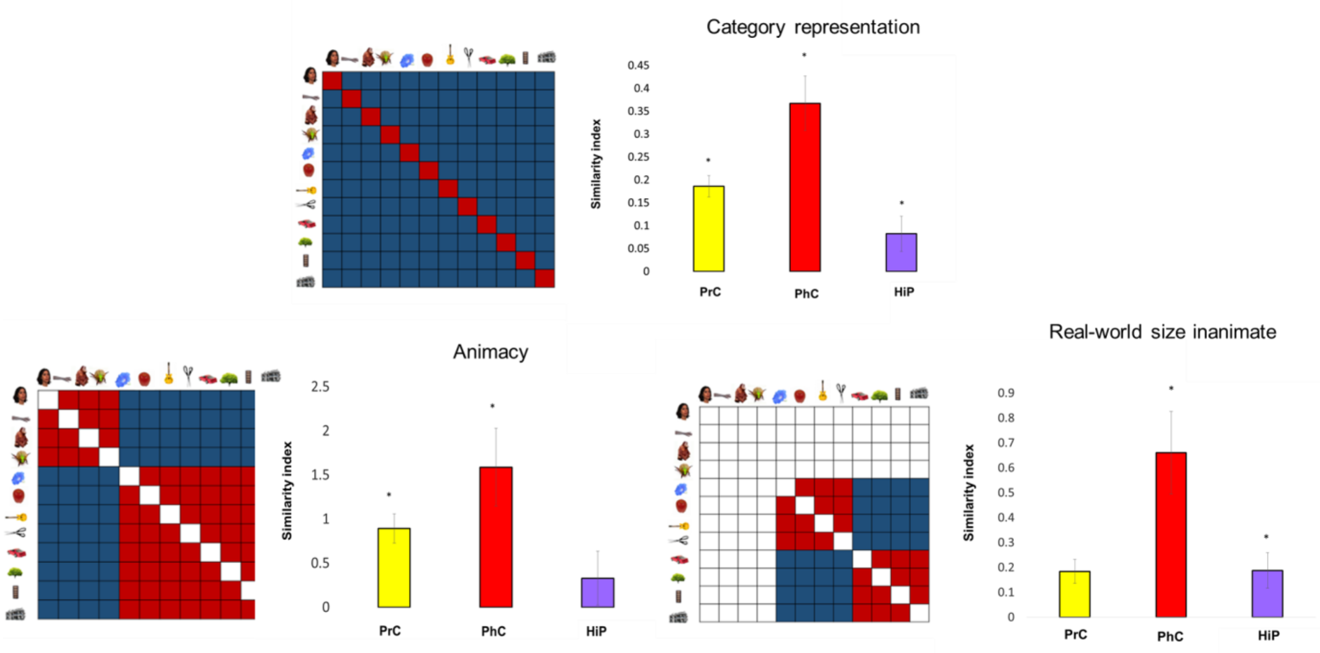

Figure 2.5: Organization of object-evoked responses in the medial temporal lobe. All Bar plots show beta fits between model of organization tested and RSM for each MTL structure. *a)* model of category representation in the MTL. All three structures show significant category representation *b)* model of animacy organization in the MTL, PhC and PrC show animacy organization PhC but not HpC *c)* model of real-world size for the inanimate domain, PhC and HpC show organization by real world size, we found no such evidence in PrC. Error bars indicate standard error of the mean *SEM*.

2.3.2.4 Visualization of representational space

In a final step we visualized the representational space for all object categories in each of the ROIs examined using hierarchical clustering (Figure 2.6). This data driven approach can be useful in that it can reveal properties that drive the organization of representations without any a priori hypotheses (Kriegeskorte et al., 2008). In PhC, the most dominant dimension of organization is that that between large inanimate objects and all other categories. In PrC, the most dominant dimension of organization is animacy. Unlike in PhC, large inanimate objects do not form a separate grouping. Finally, in HpC the most notable distinction is that between buildings and all other object categories.

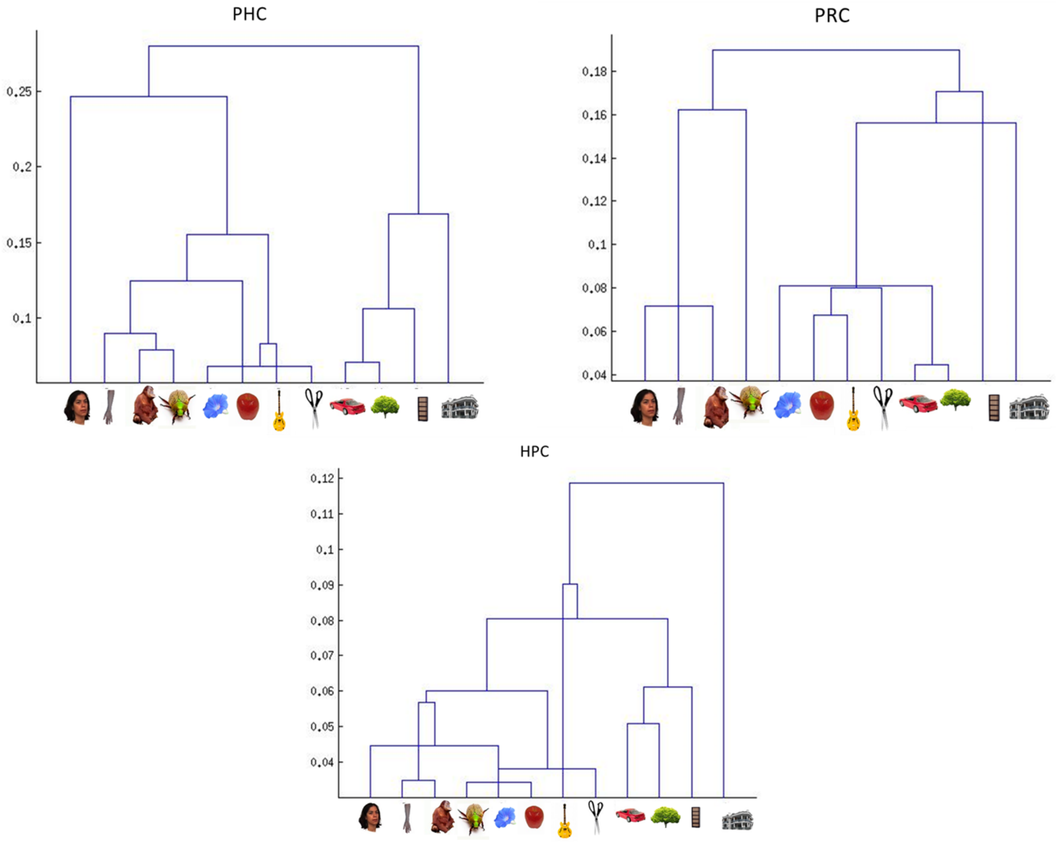

Figure 2.6: Visualization of representational space for object-evoked responses in the medial temporal lobe. Hierarchical clustering for all object categories in each MTL structure.

2.4 Discussion

In the current study we set out to explore the organization of object representations in MTL structures and to examine whether dimensions of organization prominent in upstream VVS were present in the MTL. Specifically, we asked (i) whether there is category specificity in object representations in the MTL, (ii) whether there is domain specificity along an animate-inanimate divide, (iii) whether there is specificity in representations for inanimate objects related to real-world size, and (iv) whether these dimensions of organization differ across MTL structures. We found that similar to VVS representational organization, MTL structures display sensitivity to category membership, animacy, and real-world size for inanimate objects while participants performed a recognition-memory task. Importantly, we also found differentiation across MTL structures. PhC showed organization by category, animacy, and real-world size for inanimate objects; PrC, by contrast, showed organization by category and animacy, but did not show evidence of further size-based distinctions among inanimate objects. Furthermore, there were differences in the specific types of categories represented in both structures. The HpC was largely agnostic to category membership and displayed no sensitivity to animacy, but showed sensitivity to real-world size for inanimate objects. Our findings pertaining to category specificity in the MTL replicate and extend previous findings in other task contexts. Specifically, a number of studies have reported PrC and PhC to be sensitive to object category information. In the present study we identified category specificity based on models that revealed significantly higher within- than between category similarities in patterns of activity evoked by presentation of objects. We found that PrC and PhC showed category specificity for a number of categories, while the HpC was not selective for any of the twelve categories consistent with previous studies (Huffman 2014; LaRocque 2013; Diana 2008; but see Liang et al., 2013). However we did find a trend towards category representation for buildings. In line with previous studies showing face specificity (Huffman 2014; LaRocque 2013; Liang et al., 2013), we also found a trend towards face-specificity in PrC, and in addition PrC represented items from another animate category, monkeys. PhC, by contrast, primarily represented items from large inanimate object categories, namely, buildings, furniture, and vehicles, but also faces. Some prior studies have reported decoding of faces above chance in PhC (Diana et al., 2008; Huffman & Stark 2014; Liang et al., 2013, but see Martin et al., 2013) although those using a decoding approach typically report higher decoding accuracy in PrC (Liang et al., 2013; Huffman & Stark, 2014). Specificity of response patterns for buildings and furniture has previously been reported in PhC during recognition memory judgments (Martin et al., 2013). However to our knowledge distinct representations for vehicles has not been previously explored in this structure. Taken together, our results extend previous findings that relied on a limited number of object categories that produce contiguous "blobs" across the VVS using univariate analyses. We found distinct representations for several categories in PhC and PrC, in line with findings employing representational-based analyses to explore organization across the VVS (Haxby et al., 2001; Kriegeskorte et al., 2008; Sha et al., 2015)

The sensitivity to a broader number of categories we report raises interesting questions about the principles of organization of object representations in MTL structures. Therefore, we also examined whether these structures are sensitive to two prominent dimensions that have been investigated extensively in the VVS, namely animacy and the real-world size of inanimate objects. To address this issue we examined whether models that distinguish animate from inanimate objects across all twelve categories, or between large and small objects across the eight inanimate categories, captured aspects of the representational space in each MTL structure. Importantly, we evaluated these models of animacy and size across MTL over and above any effect of category-based structure. We found that object representations in both PhC and PrC, but not the HpC, are organized along an animate/inanimate divide, whereas representations in PhC and the hippocampus, but not PrC, are organized by real-world size for inanimate objects. To our knowledge these findings provide the first evidence that principles of organization of object representations differ across all three MTL structures.

2.4.1 PrC

PrC is the MTL structure that has most extensively been linked to object processing in prior research, however to our knowledge domain-level organization has been somewhat less explored. In terms of object processing PrC has been best characterized with respect to its role in recognition memory for objects, recent work suggests that object representations in PrC also play a critical role in perceptual and semantic tasks (Barense et al., 2010; Bussey et al., 2002; Kivisaari et al., 2012, Clarke & Tyler, 2014; Bruffaerts et al., 2013; see Graham et al., 2010, for review). However, the organization of object representations that support judgements in these tasks has received only limited investigation so far. As indicated earlier, it has been reported in a number of studies that activity in PrC shows specificity for the category of faces in recognition memory and perceptual tasks (Diana et al., 2007; Martin et al., 2013, 2016; O'Neil et al. 2013, 2014). To our knowledge, domain-level organization has only been explored previously in tasks that require object naming at the basic (rather than exemplar) level. Specifically, it has been reported that PrC shows higher levels of activity when participants have to name objects that are animate as compared to objects that are inanimate **(**Moss et al., 2005**)** and there is also evidence that damage to the PrC differentially affects naming for animate objects (Wright et al., 2015). This domain-specific pattern of findings has been attributed to the fact that animate objects are distinct from inanimate objects at the level of feature statistics. Specifically, one important dimension that differs across animate and inanimate objects is the amount of feature overlap and feature distinctiveness amongst members of those domains. It has been argued that overall animate objects have more feature overlap and less distinctive features than inanimate objects (McRae et al., 1997; Devlin et al., 1998; Moss et al., 1998; Tyler et al., 2000; Tyler & Moss, 2001; McRae and Cree, 2002). In these studies feature overlap is typically defined based on listed features that can be classified as perceptual or semantic, and the level of representations tapped into by naming are at the basic level (i.e., distinguishing a horse from a zebra rather than two different horses from each other). These findings are particularly relevant for theories that view the PrC as the apex of the VVS, containing highly conjunctive object representations that allow for object disambiguation when features are highly overlapping, such as those, for example, between a horse and a zebra. Indeed, an fMRI study that employed RSA to examine object representations in PrC found that while category-level object information could be identified in more posterior VVS regions, PrC uniquely held information at the individual objects level. Specifically, patterns of activity evoked by visually presented objects were best modeled by their semantic feature overlap (Clarke & Tyler, 2014; see also Bruffaerts et al., 2013) for related findings in PrC based on written words denoting different objects). In the context of the continuous recognition memory task used in the current study, participants were required to make discriminations similar, if not more fine-grained to those required for naming an individual exemplar. Namely, the task required recognition of prior occurrence of specific exemplars, such as whether a particular building had been presented previously. Thus, although our study did not aim to test specific hypotheses about the impact of feature overlap on representational similarities, one possibility is that the animacy-related organization we report reflects differences on this dimension between the animate and inanimate objects we employed.

An alternative, not mutually exclusive account of the sensitivity of PrC to the animate inanimate distinction is that it may preferentially process animate objects due to the long range connectivity it maintains with other cortical and subcortical regions. The idea that large scale connectivity may drive differential sensitivity between stimuli of different domains, such as animate or large inanimate objects, has been investigated with regards to VVS organization in more posterior regions. Using a data-driven approach with resting-state fMRI connectivity data, Konkle & Caramazza (2016) identified three distinct resting state networks that ‘route through’ the large domain-preferring tripartite regions of VVS. Specifically, animate object preferring regions were more strongly coupled with the anterior temporal lobe, small inanimate object preferring regions were more strongly coupled with aspects of parietal cortex, and large inanimate object preferring regions were more correlated with the medial temporal lobe, as well as early visual cortex regions differentially involved in processing stimuli in the peripheral visual fields.

To our knowledge no studies have looked specifically at long range connectivity in PrC that is related to particular object domains or categories. However, several studies have examined the resting state connectivity profiles that characterize this structure (Kahn et al., 2008; Libby et al., 2012). These studies have shown that PrC shows patterns of connectivity that differ from those in PhC. For example, at the whole brain level, PrC shows distinct connectivity with other structures within the anterior temporal lobes, amygdala, and lateral orbitofrontal cortex. In terms of distinct connectivity with more posterior VVS areas, PrC shows higher connectivity with the fusiform gyrus than PhC (Libby et al., 2012). These connectivity findings have led to the suggestion the PrC is part of a cortical network, referred to as the anterior-temporal network that plays a unique functional role in memory and cognition (Ranganath & Ritchey, 2012). It has been argued that relative to a posterior-medial system of which PhC is a central component, this anterior system is preferentially involved in object recognition as well as processing the social and emotional aspects of objects and animate entities, semantic knowledge, and reward learning. Interestingly, although the Konkle et al. 2016 did not examine PrC specifically, aspects of the anterior temporal lobe were part of a resting state network connected to the posterior cortex preferring animate objects in the VVS. Although the model does not explicitly consider differences between specific object categories or domains, to the extent that the information processed in the anterior system pertains to ecologically relevant information, whether something is dangerous, this kind of processing may be more relevant to animate objects. As such, the position of PrC in this system could offer an account of PrC’s sensitivity to domain-level organization by animacy.

2.4.2 PhC

The role of the PhC in object processing during naming and recognition memory tasks has been less explored than that of PrC, including evaluating any role of feature overlap. In the memory literature, PhC has been primarily explored in terms of its role in scene recognition and in context representation in tasks of associative memory (Ranganath &

Ritchey, 2012). However, recently, it has been shown that PhC also plays a role in object recognition memory, which may be selective for certain types of objects. Specifically, it is possible to distinguish whether an object is novel or familiar from patterns of activity in PhC for some categories of objects, namely buildings, furniture, and trees (Martin et al., 2013; Martin et al., under review). In the current study we also found category specificity for buildings, furniture, and trees, and additionally for vehicles. Notably, all of these categories refer to large inanimate objects. Beyond this categorical organization, at the domain level we observed organization by animacy, and unlike PrC, PhC also showed organization by real-world size for inanimate objects. This is notable because the PPA (or parahippocampal place area), which includes a portion of posterior PhC, has also been shown to have higher levels of activity for inanimate objects, even when contrasted with shape-matched animate objects (Proklova et al., 2016). Moreover, a number of studies have demonstrated that the PPA is more active for large than for small objects (Konkle et al., 2013; Aguirre et al, 1998; Julian et al., 2016), and most similar to our findings, that patterns of activity in the PPA distinguish between large and small objects (Julian et al., 2016).

As in our discussion pertaining to PrC, it is informative to consider the long-range connectivity of PhC in relation to the category and domain level organization reported here. Resting state connectivity studies at the whole brain level have shown that PhC is differentially connected to the retrosplenial cortex (RSC), posterior cingulate, precuneus, parietal cortex, and ventromedial prefrontal cortex, the thalamus. In addition, PhC is also more strongly connected to posterior medial occipital cortex as well as early visual areas (Libby et al., 2012). In light of these resting-state connectivity findings, it has been suggested that PhC is a component of the posterior medial network, with a functional role in memory and cognition that differs from that of the anterior-temporal network that includes PrC. These findings generally align with the Konkle et al. 2016 findings that large inanimate object preferring cortex in the medial VVS is highly connected to peripheral early visual areas, as well as MTL (although not clearly specified whether it is the posterior portion of the parahippocampal gyrus, it is distinct from the anterior temporal area more highly connected to lateral animate VVS cortex). It has been argued that this network is important for representing context in episodic memory and episodic simulation, as well as in spatial navigation (Ranganath & Ritchey, 2012). One possibility is that the sensitivity of PhC to the animacy divide we report here is linked to differential processing of large inanimate objects that are important for navigation, or are more likely to serve as episodic context. Compared to animate objects, large inanimate objects often evoke a stronger sense of surrounding space (Mullally & Maguire, 2011), and when stable, can also serve as landmarks (Martin et al. under review; Janzen & van Turennout, 2004; Troiani et al., 2013). From this perspective, animacy plays a role in the organization of object representations in PhC because large inanimate objects share dimensions important for the general functions of a posterior-medial cortical system.

2.4.3 HpC

Interestingly, in our data we found that the HpC shows no sensitivity to the animacy divide and no clear-cut categorical representations of objects, although we observed a trend for the representation of buildings. Similar to PhC, the HpC was sensitive to the distinction between large and small inanimate objects. The lack of clear cut category-specific representation in our findings is in line with previous suggestions that the HpC is agnostic to the nature or content of its representations at the item level. The agnosticity of the HpC has been attributed to its unique role in pattern separation of episodes (Huffman & Stark, 2014). According to this reasoning, the result of hippocampal pattern separation is that representations in the HpC are more dissimilar to each other than those in PrC and PhC, leading to the loss of specificity in organization by category that is present in these input structures. However, the evidence for domain-level organization related to size we report here suggests that the HpC may not be entirely insensitive to content.

There is substantial evidence for a role of the HpC in scene perception and construction (Hodgetts et al., 2016; Lee et al., 2005; Barense et al., 2015; Zeidman et al., 2015; for review see Murray et al., 2017). For example, it has been demonstrated that the HpC is more active during perceptual oddity tasks for scenes than for other types of stimuli (Lee et al., 2008). Hodgetts et al. (2016) found clusters of activity in the HpC that are higher for scenes than for other stimulus categories (faces, objects) while participants performed a 1-back task, and these clusters appeared as reliably as clusters in the traditional scene-processing network (including PhC, RsC, and Transverse occipital sulcus). Based on these results, the authors suggested that the HpC should be considered as a component of the core scene processing network. Implied with this argument is the notion that the HpC is not entirely agnostic to stimulus content. More recent work by this group of researchers has provided some evidence to explain why some studies find evidence for differential involvements in scene processing and other do not (Hodgetts et al., 2017). In that fMRI study, conducted with ultra-high resolution, sensitivity to scene stimuli could be more precisely localized to a specific subfield of the HpC, namely the subiculum, with other subfields staying agnostic. It is possible that the sensitivity to real-world size of objects reported here, together with the hint for category specific representations for buildings in HpC are a result of similarities between large objects and scenes that are of particular relevance to processing in the subiculum.

2.4.4 Conclusion

Together, our findings show that stimulus properties that influence the representational structure of object information in the VVS also shape the organization of highly-specified object information in PRC, PHC, and the HpC. Moreover, they reveal that this organization differs across these MTL structures. An important direction for future research will be to test how differences in large-scale connectivity can account for these different organizational principles across structures, and how they relate to specific functional and perceptual properties of objects that differ across domains and categories. Other critical questions that deserve further investigation concern the extent to which these principles of organization are stable across different tasks and behavioral goals.

2.5 References

Aguirre, G. K., Zarahn, E., & D'esposito, M. (1998). An area within human ventral cortex sensitive to "building" stimuli: evidence and implications. *Neuron*, *21*(2), 373-383.

Barense, M. D., Henson, R. N., Lee, A. C., & Graham, K. S. (2010). Medial temporal lobe activity during complex discrimination of faces, objects, and scenes: Effects of viewpoint. *Hippocampus*, *20*(3), 389-401.

Bussey, T. J., & Saksida, L. M. (2005). Object memory and perception in the medial temporal lobe: an alternative approach. *Current opinion in neurobiology*, *15*(6), 730-737.

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

Bruffaerts, R., Dupont, P., Peeters, R., De Deyne, S., Storms, G., & Vandenberghe, R. (2013). Similarity of fMRI activity patterns in left perirhinal cortex reflects semantic similarity between words. *Journal of Neuroscience*, *33*(47), 18597-18607.

Bussey, T. J., Saksida, L. M., & Murray, E. A. (2002). Perirhinal cortex resolves feature ambiguity in complex visual discriminations. *European Journal of Neuroscience*, *15*(2), 365-374.

Capitani, E., Laiacona, M., Mahon, B. Z., & Caramazza, A. (2003). What are the facts of semantic category-specific deficits: A critical review of the clinical evidence. *Cognitive Neuropsychology*, *20*, 213-261.

Caramazza, A., & Mahon, B. Z. (2003). The organization of conceptual knowledge: the evidence from category-specific semantic deficits. *Trends in cognitive sciences*, *7*(8), 354-361.

Cohen, L., & Dehaene, S. (2004). Specialization within the ventral stream: the case for the visual word form area. *Neuroimage*, *22*(1), 466-476.

Clarke, A., & Tyler, L. K. (2014). Object-specific semantic coding in human perirhinal cortex. *The Journal of Neuroscience*, *34*(14), 4766-4775.

Cusack R, Vicente-Grabovetsky A, Mitchell DJ, Wild CJ, Auer T, Linke AC, Peelle JE (2015) Automatic analysis (aa): Efficient neuroimaging workflows and parallel processing using Matlab and XML. Frontiers in Neuroinformatics 8:90.http://dx.doi.org/10.3389/fninf.2014.00090.

Davachi, L. (2006). Item, context and relational episodic encoding in humans. *Current opinion in neurobiology*, *16*(6), 693-700.

Damasio, A. R., Damasio, H., & Van Hoesen, G. W. (1982). Prosopagnosia Anatomic basis and behavioral mechanisms. *Neurology*, *32*(4), 331-331.

De Beeck, H. P. O., Haushofer, J., & Kanwisher, N. G. (2008). Interpreting fMRI data: maps, modules and dimensions. *Nature reviews. Neuroscience*, *9*(2), 123.

De Renzi, E., Perani, D., Carlesimo, G. A., Silveri, M. C., & Fazio, F. (1994). Prosopagnosia can be associated with damage confined to the right hemisphere—an MRI and PET study and a review of the literature. *Neuropsychologia*, *32*(8), 893-902.

Devlin, J. T., Gonnerman, L. M., Andersen, E. S., & Seidenberg, M. S. (1998). Category-specific semantic deficits in focal and widespread brain damage: A computational account. *Journal of cognitive Neuroscience*, *10*(1), 77-94.

Downing, P. E., Jiang, Y., Shuman, M., & Kanwisher, N. (2001). A cortical area selective for visual processing of the human body. *Science*, *293*(5539), 2470-2473.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*(6676), 598-601.

Epstein, R., Harris, A., Stanley, D., & Kanwisher, N. (1999). The parahippocampal place area: recognition, navigation, or encoding?. *Neuron*, *23*(1), 115-125.

Eichenbaum, H., Yonelinas, A. P., & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annu. Rev. Neurosci.*, *30*, 123-152.

Epstein, R. A., & Vass, L. K. (2014). Neural systems for landmark-based wayfinding in humans. *Phil. Trans. R. Soc. B, 369*(1635), 20120533.

Frankó, E., Insausti, A. M., Artacho-Pérula, E., Insausti, R., & Chavoix, C. (2014). Identification of the human medial temporal lobe regions on magnetic resonance images. *Human brain mapping*, *35*(1), 248-256.

Fidalgo, C., & Martin, C. B. (2016). The hippocampus contributes to allocentric spatial memory through coherent scene representations. *Journal of Neuroscience*, *36*(9), 2555-2557.

Graham, K. S., Barense, M. D., & Lee, A. C. (2010). Going beyond LTM in the MTL: a synthesis of neuropsychological and neuroimaging findings on the role of the medial temporal lobe in memory and perception. *Neuropsychologia*, *48*(4), 831-853.

Grill-Spector, K., & Weiner, K. S. (2014). The functional architecture of the ventral temporal cortex and its role in categorization. *Nature Reviews Neuroscience*, *15*(8), 536-548.

Guderian, S., Dzieciol, A. M., Gadian, D. G., Jentschke, S., Doeller, C. F., Burgess, N., ... & Vargha-Khadem, F. (2015). Hippocampal volume reduction in humans predicts impaired allocentric spatial memory in virtual-reality navigation. *Journal of Neuroscience*, *35*(42), 14123-14131.

Hart, J., Berndt, R. S., & Caramazza, A. (1985). Category-specific naming deficit following cerebral infarction. *Nature*, *316*(6027), 439-440.

Hassabis, D., Kumaran, D., Vann, S. D., & Maguire, E. A. (2007). Patients with hippocampal amnesia cannot imagine new experiences. *Proceedings of the National Academy of Sciences*, *104*(5), 1726-1731.

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science, 293*(5539), 2425-2430.

Hillis, A. E., & Caramazza, A. (1991). Category-specific naming and comprehension impairment: A double dissociation. *Brain*, *114*(5), 2081-2094.

Hodgetts, C. J., Shine, J. P., Lawrence, A. D., Downing, P. E., & Graham, K. S. (2016). Evidencing a place for the hippocampus within the core scene processing network. *Human brain mapping*, *37*(11), 3779-3794.

Huffman, D. J., & Stark, C. E. (2014). Multivariate pattern analysis of the human medial temporal lobe revealed representationally categorical cortex and representationally agnostic hippocampus. *Hippocampus*, *24*(11), 1394-1403.

Humphreys, G. W., & Riddoch, M. J. (2003). A case series analysis of "category-specific" deficits of living things: The HIT account. *Cognitive Neuropsychology*, *20*(3-6), 263-306.

Janzen, G., & Van Turennout, M. (2004). Selective neural representation of objects relevant for navigation. *Nature neuroscience*, *7*(6), 673.

Johnson-Frey, A. H. (2004). The neural basis of complex tool use in humans. *Trends in Cognitive Sciences, 8*(2), 71-78.

Julian, J. B., Ryan, J., & Epstein, R. A. (2016). Coding of Object Size and Object Category in Human Visual Cortex. *Cerebral Cortex*, bhw150.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of neuroscience*, *17*(11), 4302-4311.

Kanwisher, N., & Dilks, D. D. (2013). The functional organization of the ventral visual pathway in humans. *The new visual neurosciences*, 733-748.

Konkle, T., & Oliva, A. (2012). A real-world size organization of object responses in occipitotemporal cortex. *Neuron*, *74*(6), 1114-1124.

Konkle, T., & Caramazza, A. (2013). Tripartite organization of the ventral stream by animacy and object size. *The Journal of Neuroscience*, *33*(25), 10235-10242.

Konkle, T., & Caramazza, A. (2016). The large-scale organization of object-responsive cortex is reflected in resting-state network architecture. *Cerebral Cortex*, 1-13.

Kivisaari, S. L., Tyler, L. K., Monsch, A. U., & Taylor, K. I. (2012). Medial perirhinal cortex disambiguates confusable objects. *Brain*, *135*(12), 3757-3769.

Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, *2*, 4.

Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., ... & Bandettini, P. A. (2008). Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, *60*(6), 1126-1141.

Lee, A. C., Scahill, V. L., & Graham, K. S. (2007). Activating the medial temporal lobe during oddity judgment for faces and scenes. *Cerebral Cortex*, *18*(3), 683-696.

Lee, A. C., 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.

Liang, J. C., Wagner, A. D., & Preston, A. R. (2013). Content representation in the human medial temporal lobe. *Cerebral Cortex*, *23*(1), 80-96.

Libby, L. A., Ekstrom, A. D., Ragland, J. D., & Ranganath, C. (2012). Differential connectivity of perirhinal and parahippocampal cortices within human hippocampal subregions revealed by high-resolution functional imaging. *Journal of Neuroscience*, *32*(19), 6550-6560.

Linke, A. C., Vicente-Grabovetsky, A., & Cusack, R. (2011). Stimulus-specific suppression preserves information in auditory short-term memory. *Proceedings of the National Academy of Sciences*, *108*(31), 12961-12966.

Litman, L., Awipi, T., & Davachi, L. (2009). Category-specificity in the human medial temporal lobe cortex. *Hippocampus*, *19*(3), 308-319.

Martin, A. (2007). The representation of object concepts in the brain. *Annu. Rev. Psychol.*, *58*, 25-45.

Martin, C. B., McLean, D. A., O'Neil, E. B., & Köhler, S. (2013). Distinct familiarity-based response patterns for faces and buildings in perirhinal and parahippocampal cortex. *The Journal of Neuroscience*, *33*(26), 10915-10923.

Martin, C. B., Cowell, R. A., Gribble, P. L., Wright, J., & Köhler, S. (2015). Distributed category-specific recognition-memory signals in human perirhinal cortex. *Hippocampus*.

Martin CB, Sullivan JS, Wright J, Köhler S. (2017). Recognition-memory signals for objects from different categories are graded across perirhinal and parahippocampal cortex. (under review).

McRae, K., De Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, *126*(2), 99.

McRae, K., & Cree, G. S. (2002). Factors underlying category-specific semantic deficits. *Category-specificity in brain and mind*, 211-249.

Moss, H. E., Tyler, L. K., Durrant-Peatfield, M., & Bunn, E. M. (1998). ‘Two eyes of a see-through’: Impaired and intact semantic knowledge in a case of selective deficit for living things. *Neurocase*, *4*(4-5), 291-310.

Moss, H. E., Rodd, J. M., Stamatakis, E. A., Bright, P., & Tyler, L. K. (2004). Anteromedial temporal cortex supports fine-grained differentiation among objects. *Cerebral cortex*, *15*(5), 616-627.

Mullally, S. L., Intraub, H., & Maguire, E. A. (2012). Attenuated boundary extension produces a paradoxical memory advantage in amnesic patients. *Current Biology*, *22*(4), 261-268.

Murray, E. A., Wise, S. P., & Graham, K. S. (2017). Representational specializations of the hippocampus in phylogenetic perspective. *Neuroscience Letters*.

O'Neil, E. B., Barkley, V. A., & Köhler, S. (2013). Representational demands modulate involvement of perirhinal cortex in face processing. *Hippocampus*, *23*(7), 592-605.

O'Neil, E. B., Hutchison, R. M., McLean, D. A., & Köhler, S. (2014). Resting-state fMRI reveals functional connectivity between face-selective perirhinal cortex and the fusiform face area related to face inversion. *Neuroimage*, *92*, 349-355.

Proklova, D., Kaiser, D., & Peelen, M. V. (2016). Disentangling representations of object shape and object category in human visual cortex: The animate–inanimate distinction. *Journal of cognitive neuroscience*.

Pruessner, J. C., Li, L. M., Serles, W., Pruessner, M., Collins, D. L., Kabani, N., ... & Evans, A. C. (2000). Volumetry of hippocampus and amygdala with high-resolution MRI and three-dimensional analysis software: minimizing the discrepancies between laboratories. *Cerebral cortex*, *10*(4), 433-442.

Pruessner, J. C., Köhler, S., Crane, J., Pruessner, M., Lord, C., Byrne, A., ... & Evans, A. C. (2002). Volumetry of temporopolar, perirhinal, entorhinal and parahippocampal cortex from high-resolution MR images: considering the variability of the collateral sulcus. *Cerebral Cortex*, *12*(12), 1342-1353.

Ralph, M. A. L., Lowe, C., & Rogers, T. T. (2007). Neural basis of category-specific semantic deficits for living things: evidence from semantic dementia, HSVE and a neural network model. *Brain*, *130*(4), 1127-1137.

Sacchett, C., & Humphreys, G. W. (1992). Calling a squirrel a squirrel but a canoe a wigwam: A category-specific deficit for artefactual objects and body parts. *Cognitive Neuropsychology*, *9*(1), 73-86.

Scaplen, K. M., Gulati, A. A., Heimer-McGinn, V. L., & Burwell, R. D. (2014). Objects and landmarks: hippocampal place cells respond differently to manipulations of visual cues depending on size, perspective, and experience. *Hippocampus*, *24*(11), 1287-1299.

Sha, L., Haxby, J. V., Abdi, H., Guntupalli, J. S., Oosterhof, N. N., Halchenko, Y. O., & Connolly, A. C. (2015). The animacy continuum in the human ventral vision pathway. *Journal of cognitive neuroscience*.

Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: the prospective brain. *Nature reviews. Neuroscience*, *8*(9), 657.

Troiani, V., Stigliani, A., Smith, M. E., & Epstein, R. A. (2012). Multiple object properties drive scene-selective regions. *Cerebral Cortex*, *24*(4), 883-897.

Tyler, L. K., Moss, H. E., Durrant-Peatfield, M. R., & Levy, J. P. (2000). Conceptual structure and the structure of concepts: A distributed account of category-specific deficits. *Brain and language*, *75*(2), 195-231.

Tyler, L. K., & Moss, H. E. (2001). Towards a distributed account of conceptual knowledge. *Trends in cognitive sciences*, *5*(6), 244-252.

van der Ham, I. J., Martens, M. A., Claessen, M. H., & van den Berg, E. (2017). Landmark Agnosia: Evaluating the Definition of Landmark-based Navigation Impairment. *Archives of Clinical Neuropsychology*, *32*(4), 472-482.

Walther, A., Nili, H., Ejaz, N., Alink, A., Kriegeskorte, N., & Diedrichsen, J. (2016). Reliability of dissimilarity measures for multi-voxel pattern analysis. *Neuroimage*, *137*, 188-200.

Warrington, E. K., & McCarthy, R. A. (1988). The fractionation of retrograde amnesia. *Brain and cognition*, *7*(2), 184-200.

Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain*, *107*(3), 829-853.

Wright, P., Randall, B., Clarke, A., & Tyler, L. K. (2015). The perirhinal cortex and conceptual processing: effects of feature-based statistics following damage to the anterior temporal lobes. *Neuropsychologia*, *76*, 192-207.

Chapter 3

3 Organization of recognition memory signals for objects in the medial temporal lobes

3.1 Introduction

Recognizing the objects around us, and knowing whether those objects are novel or familiar, are crucial aspects of human behavior. The importance of these abilities become starkly apparent when they are lost, for example when the ability to identify an object is lost, such as in cases of object agnosia, or when the ability to remember an object is lost, for example in cases of severe amnesia where an individual may no longer recognize the faces of family members. Much work has been done exploring how our brains enable us to identify objects (object recognition, for a review see Grill-Spector & Weiner, 2015; Kravitz et al., 2013) and remember objects (object recognition memory, for a review see: Winters et al., 2008), but a number of issues remain, particularly with regards to the relation between our brain coding for *what* an object is, and whether that object is novel or familiar.

Both object recognition and object recognition memory depend on processing in the ventral visual stream, an information processing stream running from the primary visual cortex through the inferior temporal lobe to the temporal pole (Grill-Spector & Weiner, 2015; Kravitz et al., 2013). Interestingly, there appears to be a distinction between a more general level of object recognition - i.e. recognizing that an object is a face, and recognizing that that face is a distinct entity (i.e., the face of your mother in law). This latter ability - recognizing a distinct face, is thought to be tightly linked to detecting whether that face is familiar or novel, and the representation of distinct objects, as well as recognition memory for objects, both appear to depend on the more anterior medial aspects of the VVS (Cowell et al., 2006; Murray and Bussey, 1999; Bussey et al., 2007; O'Neil et al. 2009; 2007).

A detailed account of this posterior to anterior distinction between object recognition at the category level (i.e., "a face"), object recognition at the exemplar level (i.e., a specific face), and object recognition memory comes from representational-hierarchical theory. According to this view, the specificity of an object representation increases from the more posterior aspects to the more anterior aspects of the VVS. The medial temporal lobes form the apex of the processing stream, and, accordingly, contain the most highly specific object representations, those that correspond to the unique combination of features that form distinct objects or entities. Further, these representations are the kind of representation needed to recognize an object not just for *what* it is (i.e., face of a specific individual), but for whether or not that distinct object is familiar (i.e., whether you have seen that face before). Support for this theory comes from a large body of evidence showing that one MTL structure, the perirhinal cortex (PrC), is crucial for object recognition memory. For example, the PrC is sensitive to the novelty or familiarity of objects, over longer time scales than earlier VVS areas (Fahy et al., 1993). In addition to object recognition memory, PRC is involved in object perception, particularly in cases when one is required to make fine-grained perceptual or conceptual distinctions (Clarke & Tyler, 2014; Kivisaari et al., 2012; Moss et al., 2004, Barense et al., 2007). Interestingly, there is some evidence that, like more posterior areas, distinct object representations in PrC are organized by category and domain (Liang et al., 2012; Litman et al., 2010; LaRocque et al., 2013; Huffman & Stark, 2014; Diana et al., 2008; Blumenthal et al., *under review;* Martin et al., 2013, 2016, 2017).

While PrC has primarily been the focus of object processing in the MTL, some evidence suggests the parahippocampal cortex (PhC) is also involved in object processing, although these studies have mostly been limited to looking at faces or undifferentiated object categories (Liang et al., 2012; Litman et al., 2010; LaRocque et al., 2013; Huffman & Stark, 2014; Diana et al., 2008). In the previous chapter we found categorical object representations in PhC for faces, and for four large inanimate object categories (Blumenthal et al., *under review*), and other work from our lab has shown memory signals for object categories in PhC that are partially distinct from PrC (Martin et al., 2013, 2016, 2017), specifically for buildings and furniture, whereas PrC signals are for faces and furniture. The hippocampus, on the other hand, appears to be involved primarily in associative memory (i.e., object and scene), as opposed to object memory alone; however some reports have shown distinct coding for scenes in the hippocampus (Liang et al., 2012). While we did not find category specificity for object representations in the hippocampus (although buildings was trending), we did find that object responses in the hippocampus are sensitive to the distinction between large and small inanimate objects (see Chapter 2). Therefore, whether object memory is computed in all three MTL structures, and whether this differs across structures based on their category and domain organization related to stimulus processing, has yet to be thoroughly explored.

A number of open questions remain concerning the exact relation between the organization of distinct object representations related to stimulus processing, and coding of those objects in memory signals. More specifically, it is not yet understood *how* object memory is computed in these structures, in particular whether memory is tightly linked to stimulus information, such as category or domain, or if there is a more general memory code computed across all object types. Finally, understanding how memory status, whether it be category/domain specific or more general, is implemented in MTL structures, may aid in gaining a mechanistic account of memory signaling. One way to explore this is to evaluate whether object memory is coded by average decreases or increases in neural response within a structure, or by systematic changes in the patterns of activity evoked by new or old objects.

3.1.1 Category-specific memory signals

In two studies, Martin et al. (2013, 2017) examined whether novel and familiar object representations could be distinguished from each other in PrC, PhC, and HpC, and whether this distinction differed across structures based on object categories. In both studies, they employed a study-test paradigm. Specifically, participants completed a study session outside of the scanner, during which they viewed a large number of objects from three categories. Participants then completed a test phase in the scanner. In the test phase participants were shown objects from the study list, and novel objects from the same category, and asked to respond as to whether the object was novel or familiar. A linear classifier was then trained to distinguish between brain activity patterns evoked by perceived novel and familiar stimuli within a category, and tested on a separate set of novel and familiar stimuli within that category. In the first study, the object categories employed were faces, buildings, and chairs, and in the second study, trees, planes, and chairs. It was possible to successfully decode object memory status within each category in the MTL, and this differed across MTL structures. Specifically, in PrC it was possible to classify novel from familiar faces, chairs, and planes, whereas in PhC it was possible to classify novel from familiar chairs, buildings, and trees. It was not possible to classify object memory status for any category in the hippocampus. Importantly, this classification between novel and familiar objects was only possible in a category-specific manner. In other words, a classifier trained on novel and familiar stimuli from one category, say faces, could not successfully classify stimuli from another category, such as chairs. While object memory status has been previously explored in PrC, to our knowledge there is little other work that addresses this issue in PhC; these results are novel in that they show PhC also computes object memory for some object categories. One possibility, as the authors suggest, is that the PhC is specialized for navigationally relevant object categories, objects that are generally large and stable, such as trees, and buildings, and to a lesser degree furniture. Planes, even though they are large, due to their frequent movement, do not serve as landmarks, and object memory signals for this category were found in PrC and not PhC. It is also notable that there appears to be a gradient and not a sharp distinction between structures, with memory signals for furniture being decodable in both structures. While these results suggest interesting differences in object memory representation between structures, exploration with a larger number of categories that differ across domains could help to further map out these distinctions. Additionally, the finding that decoding was category specific opens up the question as to whether this is a general property of how these structures are organized and contribute to memory. Therefore, it is important to see whether memory signals are category specific for a larger number of categories, as well as domain specific, and whether this is a general property in the organization of these regions, or varies by task-space.

3.1.2 Coding of memory signals

A wealth of research has shown that PrC plays a critical role in object recognition memory in humans, non-human primates, and rodents (Brown et al., 2001; Murray et al. 1999, 2007; Winters et al. 2004). Much of this work suggests that PrC codes for object novelty or familiarity through repetition suppression - or a decrease in neural response upon repeated exposure (or a decrease in response for a previously familiar stimuli). Recordings from monkey PrC have shown that subpopulations of neurons in PrC show decreases in activity for repeated stimuli on a continuous recognition memory task (Fahry et al., 1993), as well as on delayed-match to sample or non-match to sample tasks, where some objects incidentally become familiar across trials (Li et al., 1993; Miller et al., 1993). Interestingly, PrC neurons have been reported that show a rapid single-trial decrease in firing (Fahry et al., 1993). In addition to this rapid responsiveness, neurons have been reported that continue to show decreases in response even after a number of intervening stimuli, and after delays as long as 24 hours (Fahry et al., 1993; Xiang & Brown, 1998). This sensitivity to object repetition at longer time-scales, such as 20+ trials of intervening stimuli, or 24 hour delays between recording sessions, differentiates repetition responses in PrC from more posterior VVS areas that show repetition only over short delays of a few seconds (Miller et al., 1993). Xiang and Brown (1998) categorized these PrC neuronal responses into 3 types: neurons that respond to recency of stimuli presentation (recency neurons), neurons that respond to familiarity (repeated exposure prior to testing), and novelty neurons. While both recency and familiarity neurons showed distinct decreasing response profiles, interestingly, in some cases, when stimuli became *highly* familiar over 20 presentations, a subpopulation of neurons actually showed differential *increases* in response. In fact, while decreases, or repetition suppression, have been most widely reported, other studies have reported increases. Miller et al. (1993) reported that some PrC neurons showed increased firing after hundreds of presentations of the same items in a DMS task, and Li et al. (1993) showed that on a DMS task, a small population of neurons (relative to those that showed decreases in response) showed increases, in conjunction with overall increases in baseline firing. These electrophysiological studies show compelling evidence that PrC neurons are sensitive to object memory status, potentially for long-term memory, however a precise understanding of how these memory differences are coded (by increases, decreases, or a combination) remains an open question.

A further challenge is linking these findings to the human fMRI literature. Similar to electrophysiological findings, changes in overall levels of neural activity as measured with BOLD have been shown to differentiate novel from repeated stimuli. In a meta-analysis of 48 fMRI studies, Kim (2011) showed that the bilateral medial temporal lobe reliably shows decreased activity when contrasting activity during hits (correctly identified old objects) with correct rejections (correctly identified new responses). Importantly, robust repetition suppression for object stimuli, as measured by BOLD decreases, has been reported in studies focusing on PrC (Heusser et al., 2013; Duke et al., 2017; Wang et al., 2014), but has also been reported in both PhC and PrC (Yassa & Stark, 2008; Gonsalves et al., 2005) in whole brain analyses. In PrC, these decreases in activity have been found to parametrically track repetition (Yassa & Stark, 2008; Duke et al., 2017), similar to the recency responses described in the electrophysiology literature. In PhC and PrC, decreases in response to old stimuli during retrieval track the strength of recognition memory, and in PrC decreases have also been shown to relate to implicit conceptual priming (Heusser et al., 2013; Wang et al., 2014). However, much like in the animal electrophysiology literature, the picture is more complicated - increases in activity related to repetition or familiarity have also been reported. Yassa and Stark (2008) showed that in a continuous recognition memory task, an area in anterior PRC shows decreases related to repetition, while an area in posterior PrC shows increases in BOLD response that track memory strength. Finally, Duke et al. (2017) showed that while recency is tracked by PrC suppression, the same area of PrC shows parametric increases with lifetime familiarity for an object.

The above results are based on a univariate analysis approach to evaluating how PrC codes for memory. An alternative approach is to ask whether pattern based distinctions between novel and familiar stimuli are driven by increases, decreases, or a combination of both. This is particularly relevant in mvpa approaches given that a classifier can leverage any of these changes to distinguish novel from familiar object evoked responses. To explore this, Martin et al. (2016) conducted a follow up study using the same data that showed category-based memory signals in PrC for faces, houses and chairs (Martin et al., 2013). They addressed this question by selecting voxels in PrC that showed the greatest decreases, increases, or absolute changes between novel and familiar face presentations, and running the classification analysis using these voxels. All three forms of feature selection led to above chance classification between novel and familiar faces. However, interestingly, only classification accuracy using a combination of voxels that showed increases and decreases correlated with recognition memory performance across participants, with better classification correlating with better discrimination between novel and familiar faces. These findings suggest that memory signals are coded by both increases and decreases, however as the results are specific to faces in PrC, further research is required to understand whether this is true for a number of categories in different MTL structures.

3.1.3 Goals of the current study

In the current study, we replicated and extended Martin et al. (2013, 2016, 2017). We examined whether category-specific memory signals for a broad number of categories were present in MTL structures, and whether this category-specificity in the memory signal differed across structures. Additionally, based on our findings that object responses are organized by the domains animacy and real-world size during stimulus processing, we explored whether there were domain-specific memory signals, and whether this organization differs across MTL structures. We hypothesized that structures that represented a certain category in general would also show category-specificity in the memory signal, and further, structures that showed a domain organization in general would also show domain-specific memory signals. In Chapter 2, we evaluated object evoked response patterns regardless of memory status, and found that PhC distinctly represented faces, buildings, furniture, trees, and vehicles, whereas PrC represented monkeys, and at a lower statistical threshold faces, and the HpC showed no category specificity, other than a trend for representing buildings. Based on these findings, we predicted that we would also find category-specificity in memory signals for these categories in each structure. Expanding on the idea that memory signals are stimulus specific, or at least category specific in nature in the MTL, we also asked if memory signals carried domain specific information, and whether this varied by the domain organization in each MTL structure. Specifically, based on our domain-level findings in Chapter 2, we hypothesized that it might be possible to distinguish novel from familiar patterns of activity within the animate or inanimate domains in PrC and PhC, and novel from familiar patterns within large or within small inanimate objects in PhC and the HpC. We took advantage of the continuous recognition memory task data used in Chapter 2 to explore these questions. Importantly, unlike in study-test paradigms where novel and familiar stimuli in the test phase are necessarily different subsets of stimuli, in a continuous recognition task stimuli are identical - the only difference is the memory status of that item. This design optimally ensures that distinctions between novel and familiar stimuli are due to memory status, and not to any perceptual or image-based differences between stimuli sets. We utilized the advantage of this design to investigate whether or not memory signals reflected category or domain, as well as to explore whether or not these distinctions were coded by repetition suppression, average increases in activity, or by changes across patterns of activity. Further, we asked whether pattern based changes were the result of average increases, decreases, or the pattern across increases and decreases. In the first set of analyses, we isolated memory signals for each category in each MTL structure, and asked whether the representational space of these memory signals were organized by category and domain. In the second set of analyses, we asked whether it was possible to distinguish novel from familiar object evoked responses within a category or within a domain in each structure. In the final set of analyses, we asked whether there was repetition suppression for each category in each ROI, and whether there was a change above and beyond repetition suppression in the patterns of activity for novel and familiar objects.

3.2 Materials and methods

The data used in the current study are the same data used for a separate study, “Animacy and real-world size shape object representations in the medial temporal lobes” which was described in Chapter 2 of the current thesis. Therefore, please refer to the Chapter 2 methods section, for information on participants, stimuli, experimental procedures, and image acquisition, which are identical. Importantly, in the current study the data were modeled differently. In general, multivariate analyses employed in the current study were highly similar to those described in Chapter 2, but procedures differed in some details in order to explore questions related to the memory status of object representations. Therefore, new methods sections are outlined below.

3.2.1 Neuroimaging analyses

3.2.1.1 Pre-processing and modeling

fMRI data were analyzed using SPM8 (Welcome Institute of Cognitive Neurology; http://www.fil.ion.ucl.ac.uk/spm/software/spm8/), employing an analysis pipeline as implemented in the automatic analysis system (aa) (www.github.com/rhodricusack/automaticanalysis), (Cusack et al., 2015). Functional data were motion corrected and high pass-filtered to remove low frequency noise (drift); slice-time correction was not implemented due to the use of a multiband sequence. Four dummy scans at the start of each session were discarded to allow for T1 relaxation. Functional Images were then co-registered with participant-specific anatomical images. Following co-registration an MVPA pipeline was initiated in aa, co-registered images were kept in native space for each participant, and no spatial smoothing was applied in order to preserve high-spatial resolution. Functional data were convolved using a canonical hemodynamic response function. In order to explore memory signals, only trials in which the participant responded correctly on the 1st exemplar presentation (novel) and correctly on the 2nd presentation of the exact same exemplar (familiar) were modeled and included in the analysis. With this procedure, any comparisons between new and old stimuli differ *only* in memory status, since all images are identical. Regressors comprised the onsets and durations of these stimuli (onset + 1200ms). While performance was high overall, for some participants, on some runs there were no trials in a certain category condition (e.g., “familiar faces”). In order to have enough power to model conditions with this performance criteria, modeling was collapsed across odd (1, 3, 5, 7) and even (2, 4, 6) runs. Therefore, regressors were modeled at the category level for novel and old stimuli for odd runs and for even runs, for a total of 24 regressors per odd runs (e.g., “novel faces”, “old faces”, etc.) and 24 regressors for even runs (e.g., “novel faces”, “old faces”, etc.).With this modeling, beta estimates for each category were derived based on a maximum of 16 exemplars for novel, and 16 exemplars for familiar for odd runs, and 12 exemplars for novel and 12 exemplars for familiar for even runs (assuming the person was correct on every trial for that category), but differed based on subject’s performance. Regressors of no interest included 6 motion regressors. Beta estimates derived from these models were either subjected to a univariate contrast, after which the contrast values were used as input to the multivariate analyses, in the first set of analyses described, or were directly used as input for multivariate analyses, for the 2nd set of analyses described. Due to the high level of control, i.e., only modeling stimuli in which participants were correct on both exemplar presentations, we necessarily had less trials to model than in previous analyses (those described in chapter 1). Therefore, we later ran the identical set of analyses using data modeled on all correct novel and familiar stimuli, in order to ensure that the results were not due to power issues. We found that the results replicated, and therefore we report only the results from our targeted analysis based on trials in which participants correctly identified the novel and familiar exemplar. Medial temporal lobe ROIs were demarcated manually for each participant on the high-resolution structural images in native space, using the anatomical protocols published by Pruessner et al. (2000; 2002) with adjustments to the posterior border of PhC as specified by Franko et al. (2014).

3.2.1.2 Feature-selection: Isolating memory signals

An important goal of the current study was to isolate brain activity patterns that corresponded to the memory component of the signal. With this in mind, we designed a continuous recognition memory task. Most traditional memory studies that employ a study-test paradigm, and evaluate activity at test for the familiar (i.e., studied) items, compared to a different set of novel items. In contrast, in a continuous recognition memory task, novel and familiar stimuli are identical. The benefit of this approach is that any differences between novel and familiar stimuli must be due to their memory status alone, since there are no visual stimuli differences. We reasoned that in terms of brain activity, any difference in response between correctly identified novel (1st presentation) and correctly identified familiar (2nd presentation) trials of the same exemplar would correspond to the "memory" component of the response. Therefore, in our first set of analyses, to isolate the memory signal for each category, we conducted a univariate contrast between novel and familiar objects for each category, on a subject by subject basis, for odd runs and for even runs. This contrast was conducting using all voxels in each predefined ROI (PrC, PhC, and HpC), and resulted in 12 "memory signals" or contrast values across new and old objects for each category for odd runs, and for each category for even runs. Four different feature selection approaches were run using these contrast values. In the first schema, contrast values for all voxels were included in the analysis. In the second approach, in each ROI we selected the 20% of voxels with contrast values that showed the largest absolute difference between novel and familiar stimuli, in other words, voxels that showed the largest change between conditions regardless of direction of change. This was motivated by previous work in our lab showing that both increases and decreases in activity across voxels in PhC and PrC correlates with recognition memory behavior (Martin et al., 2013, 2016). In the third selection schema, motivated by the large literature on repetition suppression, or decreases in activity levels in response to repetition, we included the 20% of voxels in each ROI that showed the largest decrease across first and 2nd presentation. For the fourth feature selection schema, based on studies that show increases with repetition or familiarity, we included the top 20% of voxels that showed the greatest increases between novel and familiar stimuli. These methods of feature selection allowed us to explore whether increases, decreases, or a combination of both characterize memory signal in MTL structures.

3.2.1.3 Representational similarity analysis: Exploring organization of memory signals

After extracting patterns of activity across voxels in each ROI that corresponded to the memory signal in each category, we used these category-based memory patterns as features (or input) to our multivariate analyses (Kriegeskorte et al., 2008). Representational similarity analyses were identical to those described in Chapter 2, with the important distinction that these analyses were now conducted on the memory signals, or difference in activity between novel and familiar, as opposed to modeled on the activity in response to stimulus presentation regardless of memory status. Additionally, unlike in Chapter 2 where we were able to model categories in each of the 7 runs and compute between run correlations, here we modeled stimuli in odd and even runs, and computed correlations for category memory signals between the odd and even runs. Importantly, similar to the between run approach used previously, computing correlations across odd and even runs ensured that comparisons were not made across regressors with close temporal proximity (Linke et al., 2011).

To explore the memory-signal representational space in each ROI, for each participant, we computed the Pearson’s correlation for each category memory signal compared to itself, as well as to each other category memory signal. This resulted in a 12 x 12 representational similarity matrix (RSM) of the memory signal for each participant, for each ROI, with within category similarity values (between odd and even runs) on the diagonal, and between category information (between odd and even runs) on the off diagonal. RSM’s were then averaged across participants, resulting in a final group similarity matrix for each ROI (see Figure 3.1). For visualization purposes, group RSMs were ordered in the following way: animate objects, small inanimate objects, and large inanimate objects. We then tested whether the memory signal itself reflected the categorical and domain level organization that was present when modeling brain activity patterns in response to stimuli regardless of memory status (see Chapter 2). As in our previous study, we asked whether the representational space was modulated by category, animacy, and size within inanimate objects. To do this we created linear models (predefined contrasts) specifying which RSM correlation values were to be subjected to a t-test (see Figure 3.2). These analyses were performed on data in single-subject RSMs, with the group statistics calculated from the average results.

We first asked whether there was evidence of a category-level organization of the memory signals in each ROI. To test for this, identical to our approach in Chapter 2, we defined a contrast of category representation (see Chapter 2 Figure 2.5), in other words a linear model where all within category (diagonal) memory signal patterns were more hypothesized to be more highly correlated than between category memory signal (off diagonal) patterns. In the initial analysis, we tested an omnibus contrast (i.e., model) that probed for the presence of any category-specific information in each ROI. We then tested for information relating to each of the 12 categories individually. Specifically, we tested whether the patterns of the difference signal, across voxels were more similar within each category than between that category and the 11 other categories. Importantly, this analysis allows us to evaluate whether memory signals are organized categorically.

In our second set of analyses, we asked whether or not the animate versus inanimate object distinction that we found to shape the organization of object representations in the MTL in our previous set of analyses was also present in the object-category memory signals. Specifically, we tested whether the memory signals for each category, when compared to other categories within their domain, still showed domain level similarities in correlations (i.e., whether memory signals within animate objects and within inanimate objects were more similar to each other than between domains) (see Chapter 2, Figure 2.5). We hypothesized that the same areas that showed sensitivity to animacy in our previous analyses (PhC and PrC) while viewing stimuli without isolating a memory component, may also show domain organization in the memory signal. Importantly, in these analyses we removed the diagonal from our model in order to discard any possible influence of within-category memory similarities.

In our third and final set of analyses, we asked whether real world size is an organizing dimension for memory signals within the domain of inanimate objects in MTL.

Specifically, we hypothesized that PhC and HpC memory signals might be sensitive to size, as these areas were sensitive to real-world size for stimuli irrespective of memory status. Here, we divided inanimate objects into groups of small or large objects, large objects included trees, furniture, vehicles and buildings, and small objects including fruit, flowers, musical instruments, and tools. This analysis was identical to the previous one, except that within versus between similarities were computed across all categories of large or small inanimate objects (see Chapter 2, Figure 2.5). As in the analyses on animacy described above, we did not include the diagonal when testing this model.

3.2.1.4 Representational similarity analysis: Exploring perceived memory status

In our first set of analyses described above, we isolated object category memory signals and explored the representational space of those signals. In particular, we asked whether the memory signals themselves reflected stimulus category and/or domain information. In the second set of analyses we explored the same question - whether object memory in these regions is organized categorically or by domain, employing a different approach. In this approach we did not conduct a contrast between old and new category responses, but instead asked if we could distinguish between old and new object presentations at the category or domain level. To do this we computed Pearson's correlations between patterns of activity evoked across all voxels for all novel categories and all familiar categories (using correct responses on both exemplar presentations as previously described) for each subject in each ROI, resulting in a 24 X 24 matrix, with novel objects consisting of the first 12 cells (in order by animacy and size for inanimacy) and familiar objects for the last 12 cells. We then asked, for each category, whether you could distinguish perceived novel from perceived familiar conditions, by hypothesizing that patterns of activity for novel objects for that category would be more similar to each other, than to patterns of activity for familiar objects in the same category, and vice versa. Finally, we asked whether we can distinguish novel from familiar patterns of activity within our domains of interest. Specifically, we asked whether patterns evoked by novel animate objects were more similar to each other than to patterns evoked by familiar animate objects, whether patterns evoked by novel inanimate objects were more similar to each other than patterns evoked by familiar inanimate objects, and lastly whether patterns evoked by novel large and small objects were more similar within the novel presentations than old, and vice versa. These hypotheses were tested using a linear models (see Figure 3.3) that were then regressed onto the RSM, in an identical approach to the analysis described in the previous chapter, but now comparing novel and old patterns of activity. As in the previous analyses, we removed the diagonal in our domain level models. Lastly, we asked whether or not you could distinguish between novel and familiar object presentations across all object categories (see Figure 3.3).

3.2.1.5 Across-category memory signal: repetition suppression

To foreshadow the results of the set of analyses just described, we found that for all ROI's it was not possible to distinguish between novel and familiar object responses within any category, or within any domain. However, we did find that it was possible to distinguish novel from familiar responses across all object categories. One possible explanation of this general effect is that the pattern distinction is driven by a repetition effect, or a change in average signal response across all voxels within an ROI. Previous work has shown that pattern similarity is affected by average changes in the signal across conditions, specifically, when using Pearson's correlations as a similarity measure; higher levels of activity can increase pattern similarity by shifting the origin of the patterns closer together (Walther et al., 2015). Therefore, we explored whether there was evidence for repetition suppression in each of our ROI's, and whether repetition suppression, defined as an average decrease in activity upon second repetition, could account for the differences in similarity between new and old patterns in their entirety. To explore this issues, we used Euclidean distance as a similarity measure in combination with a per-voxel pattern demeaning method within our conditions of interest (novel and familiar), which allowed us to effectively remove any effect of average activity levels on pattern similarity (Walther et al., 2016). First, to evaluate whether there was repetition suppression, we compared the Euclidean distance for each novel and familiar category relative to 0 (baseline) (Figure 3.4), and computed a t-test for novel response vs. familiar response across all object categories, and for each object category compared to itself, in each ROI.. After evaluating whether each ROI shows significant repetition suppression, we then performed pattern demeaning to remove the average change in activity across novel and familiar conditions. Specifically, within the novel condition and within the familiar condition, we subtracted the mean activity level across all categories on a voxel by voxel basis before computing a Euclidean dissimilarity matrix, which we then tested with our linear model that specifies a distinction between novel and familiar responses across all object categories (Figure 3.3).

3.3 Results

3.3.1 Behavioral

A full table reporting *d'* and reaction times for all correct rejections (novel trials) and hits (familiar trials) can be seen in Table 2.1 of chapter 1.

3.3.2 fMRI results

3.3.2.1 Representational similarity analysis: Exploring memory-signal organization

In our first set of analyses, we isolated the memory signal for each category by conducting a univariate contrast between novel and familiar objects in that category. We then used these differences, or contrast values, as input to our representational similarity pipeline. Specifically, we computed correlations between memory signals within and between all categories in each participant for each ROI, and averaged these memory signal RSMs across participants. Figure 3.1 shows the average memory signal RSM, or representational similarity matrix for each MTL ROI examined (PrC, PhC, and HpC).

The RSMs displayed here and subsequent analyses reported include all voxels in each ROI, but analyses were also conducted with initial feature selection, using the 20% of voxels showing the greatest change between novel and familiar trials, the greatest decrease between novel and familiar trials, or the greatest increase between novel and familiar trials. However since results did not differ, data are reported for only the first feature-selection approach in which we included all voxels. We then asked whether there was an effect of category in each ROI, by testing a model of higher within-category than between-category similarity. We found no evidence of category representation in any ROI (PhC = $t(12) = 1.02$, p=0.32, PrC =$t(12)$ =0.64, $p = 0.53$, HpC = $t(12) = -0.36$, $p = 0.72$). Individual category tests, in which we tested within category compared to between category similarity (12 tests) also revealed no significant effect for any ROI (see Table 3.1). Lastly, we tested whether memory signal representational space was organized by domain. We tested a model of animacy and a model of real-world size for inanimate objects (see Figure 2.5, Chapter 2). Results were not significant in any ROI for either animacy organization (PhC = $t(12) = 0.88$, p=0.39, PrC =$t(12) = 0.48$, $p = 0.63$, HpC = $t(12) = 0.63$, $p = 0.54$), or real-world size organization for inanimate objects (PhC = $t(12) = -0.73$, $p = 0.48$, PrC =$t(12) = 0.46$, $p = 0.65$, HpC = $t(12) = -1.46$, $p = 0.17$) .

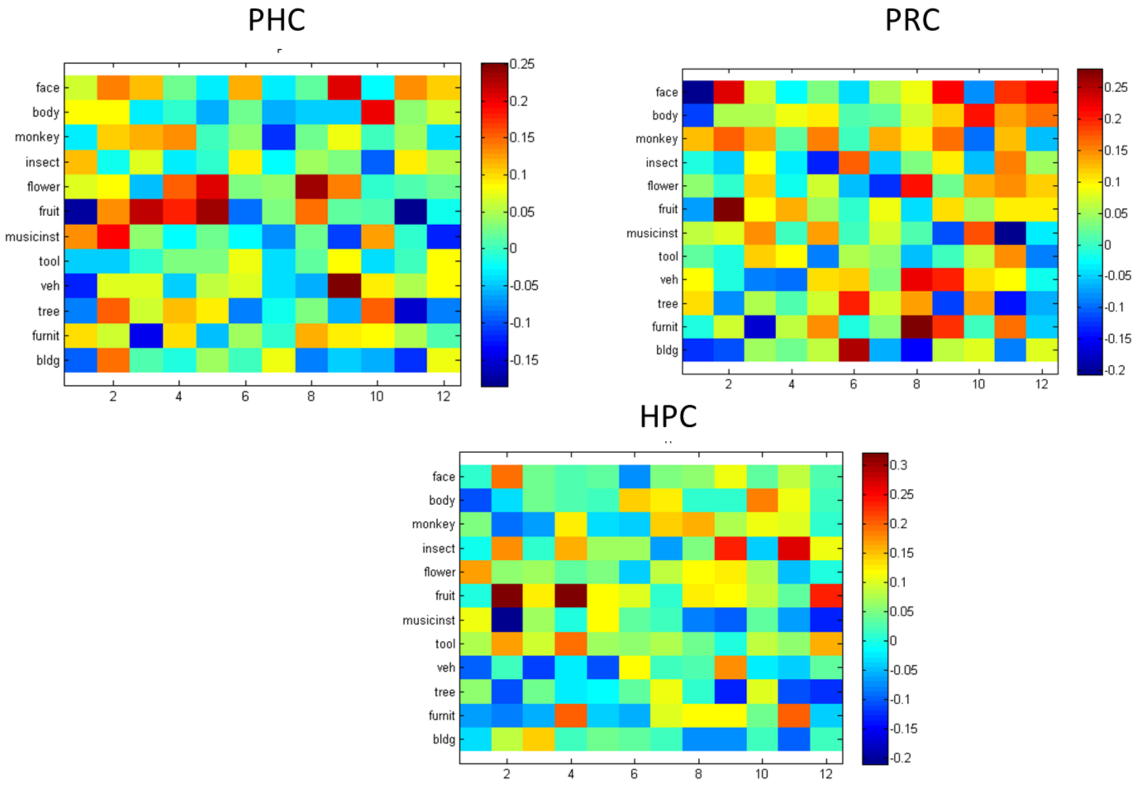

Figure 3.1: Recognition-memory signal representational similarity analysis for MTL structures. Pearson's correlations between object memory signals (univariate difference between novel and familiar stimuli) for each of the 12 categories compared to each other, warmer colors indicate a higher similarity.

	PhC				**PrC**				**HpC**		
	beta	T	p		beta	T	p		beta	T	p
face	0.010	-0.12	.90		0.02	0.38	0.71		-0.08	-1.29	0.22
body	-0.08	-1.84	0.09		-0.006	-0.09	0.93		-0.09	-1.3	-0.19
monkey	-0.08	-2.04	0.07		0.002	0.04	0.96		-0.02	-0.40	0.70
insect	0.021	0.32	0.75		-0.05	-1.41	0.27		-0.04	-0.73	0.48
fruit	0.003	-0.05	0.96		-0.01	-0.26	0.79		-0.07	-1.02	0.33
flower	-0.00	-0.02	0.98		0.07	1.38	0.19		-0.07	-1.32	0.21
music instrument	0.015	0.23	0.82		-0.06	-1.32	0.21		-0.04	-0.53	0.60
tool	0.02	0.35	0.73		0.03	0.70	0.49		-0.05	-0.84	0.41
vehicle	0.14	1.56	0.14		0.17	1.94	0.07		0.16	1.67	0.12
tree	0.05	0.89	0.38		0.13	2.22	0.05		0.05	2.00	0.07
furniture	0.09	2.28	0.04		0.04	0.83	0.42		0.08	1.42	0.16
building	-0.02	-0.22	0.48		0.12	0.46	0.65		-0.33	-1.46	0.17

Table 3.1: Results from category specific linear contrast tests on the recognition-memory signal in each MTL structure tested. The beta indicates the fit between a within-versus between linear contrast on the memory signal RSM, results are Bonferroni corrected for multiple comparisons (12 category tests for each ROI).

3.3.2.2 Representational similarity analysis: Exploring perceived memory status

In our second set of analyses, we addressed the same question, i.e., whether or not the memory component of the signal in MTL contained category or domain information, using a different approach. Specifically, we tested whether it was possible to distinguish novel and familiar patterns of activity within categories or within domains. For each category, we tested a model that novel object evoked patterns from that category would have more similarity to each other as compared to familiar object evoked patterns, and vice-versa (Figure 3.2). We were not able to distinguish perceived novel from familiar evoked responses for any of the 12 categories in PhC, PrC, or HpC. Next we asked whether it was possible to distinguish novel from familiar object evoked responses within domains (e.g., whether novel animate objects evoke more similar patterns of activity to each other than to old animate objects) (Figure 3.2). We also found no evidence of a distinction in activity patterns by perceived memory status for animate object categories (PhC = *t*(12) = -1.91, *p* = .08, PrC =*t*(12) = -2.09, *p* = .06, HpC = *t*(12) = 0.07, *p* = .94), nor for inanimate object categories (PhC *t*(12) = -0.33, *p* = 0.75, PrC =*t*(12) = 1.60, *p* = .13, HpC *t*(12) = 0/02, *p* = .98), or for large inanimate object categories (PhC *t*(12) = -0.66, *p* = .52, PrC *t*(12) = 0.21, *p* = .83, HpC *t*(12) = -1.85, *p* = .09), nor small inanimate object categories (PhC *t*(12) = 0.02, *p*= .97, PrC *t*(12) = 0.42, *p* = .67, HpC *t*(12) = 1.3, *p* = 0.20). As a final analysis, we asked whether memory status overall (regardless of object category) could be detected in differing patterns of activity in each MTL region (Figure 3.3). Interestingly, we found that in each ROI, we could distinguish perceived novel and familiar objects across all object categories (i.e., in analyses that were blind to category; PhC t(12) = 2.93, *p* = 0.01, PrC *t*(12) = 2.80, *p* = 0.01, HpC =*t*(12) = 2.29, *p* = 0.03).

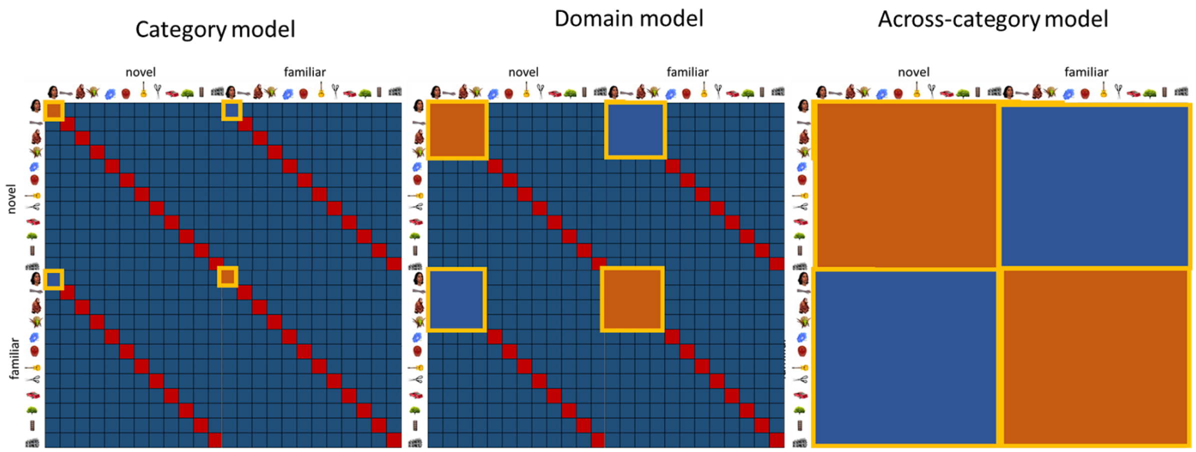

Figure 3.2: Example of linear models used to test within category, within domain, and across category distinctions between perceived novel and familiar object stimuli. The first 12 columns and rows indicate data for perceived novel trials, and the last 12 columns indicate data for perceived familiar trials. Warmer colors indicate a higher hypothesized similarity. Models tested are indicated in the gold outlined boxes, using faces as an example for within-category, and animate objects as an example for within-domain tests. Note that in each category and domain test the other cells (indicated by the red diagonal and navy cells outside of the gold outlined example) were not included in each model.

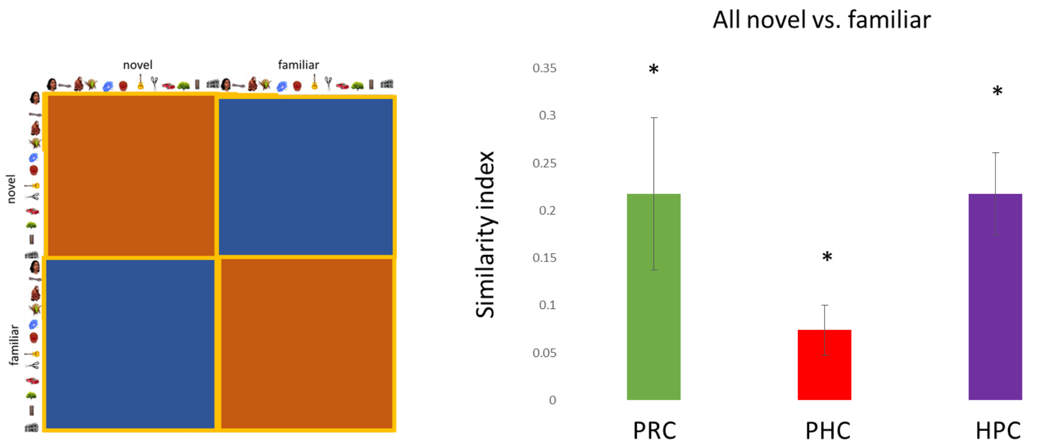

Figure 3.3. **Across category distinction between perceived novel and familiar object stimuli. Example linear model used to test the hypothesis that you can distinguish perceived novel from familiar object stimuli across all object categories, because novel object representations are more similar to each other than to familiar object representations, and vice versa (left). Results from the across-category linear contrast test for each MTL structure, the y axis shows a similarity index, which is the beta value of the fit between the linear model and the RSM in each structure. Memory status across all object categories can be distinguished in each MTL structure.**

3.3.2.3 Across-category memory signal: repetition suppression?

We examined whether or not a general repetition suppression effect, or decrease in average activity across all voxels in our regions of interest, could account for the distinction between patterns related to all perceived novel versus familiar object categories. We found significant repetition suppression across all categories in each ROI (PhC = $t(12) = 3.42$, $p = 0.001$, PrC =$t(12) = 4.39$, $p = 0.0005$, HpC = $t(12) = 2.50$, $p = 0.01$). We then tested each individual category within each ROI to see whether there was significant repetition suppression, using Bonferroni correction for multiple comparisons (12 tests in each ROI). In HpC we found significant repetition suppression for furniture ($t(12) = 3.27$, p =0.03), in PhC we found repetition suppression for insects ($t(12) = 3.50$, p = 0.02, musical instruments ($t(12) = 3.43$, $p = 0.02$, and furniture ($t(12) = 4.39$, $p = 0.0005$, and in PrC we found significant repetition suppression for bodies $t(12)$ =3.28, $p = 0.03$, and furniture $t(12)$ =3.68, $p = 0.01$) (Figure 3.4). We then removed the repetition suppression effect by subtracting mean activity in each condition (novel and familiar) on a per voxel basis, and then examining the Euclidean distance between patterns for novel versus familiar items. After removal, there was still a significant difference between patterns for all novel versus all familiar items across all object categories in each MTL structure, and, interestingly, patterns of activity for novel objects were more dissimilar amongst themselves than patterns of activity for old objects within themselves (PhC = $t(12) = -2.41$, p =0.01, PrC =$t(12) = -2.68$, $p = 0.007$, HpC = $t(12) = -2.87$, p = 0.004) (Figure 3.5)**.**

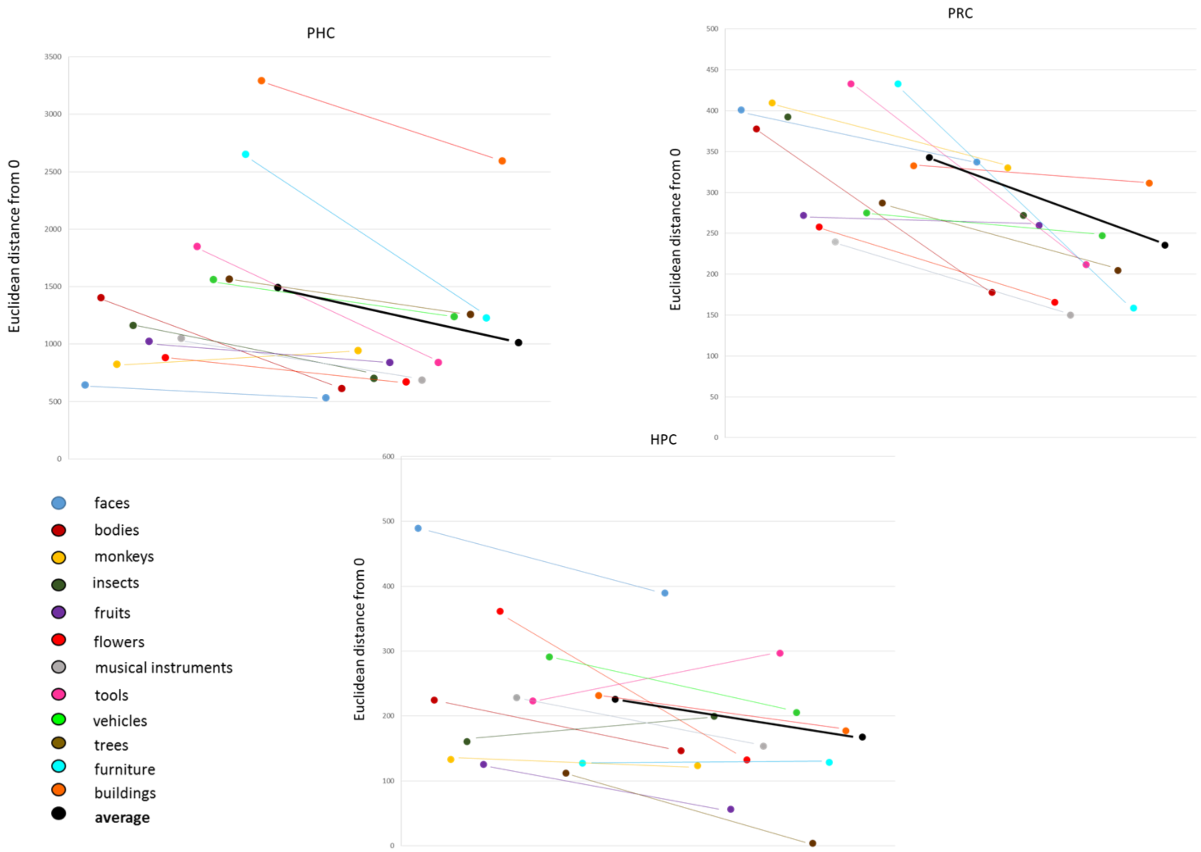

Figure 3.4: Repetition suppression in MTL structures. Repetition suppression is measured by Euclidean distance from 0. Each category is plotted in a different color, with the novel activity as the first data point, connected by a line to the familiar category data point. The average between all categories is shown by the black line. Repetition suppression, or an average decrease in BOLD activity was significant across all object categories in each MTL structure.

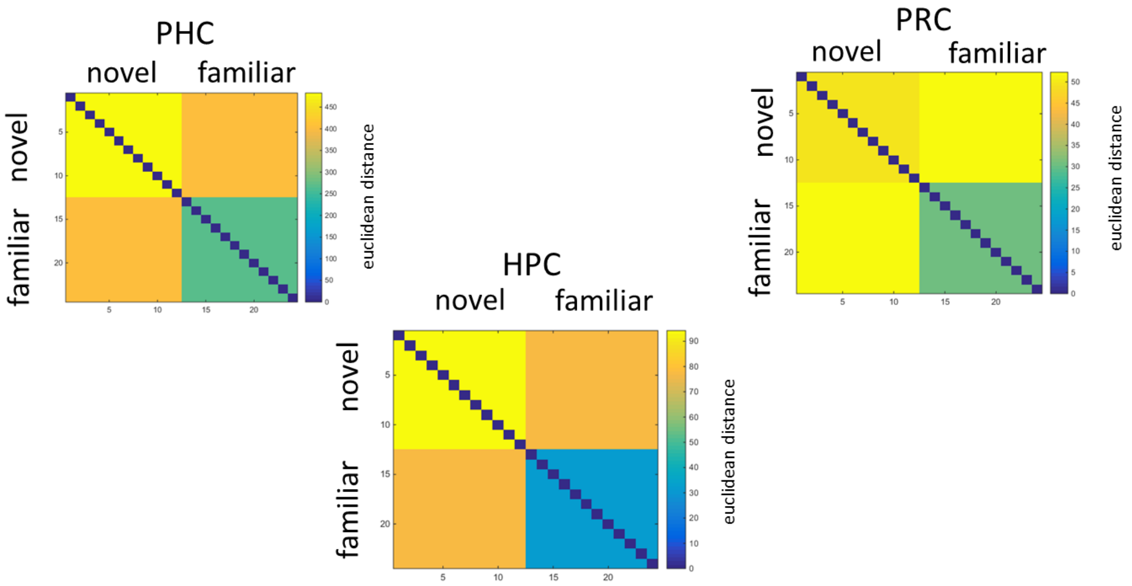

Figure 3.5: Representational dissimilarity matrices averaged across category for novel and familiar object stimuli, after pattern-based demeaning. There is a significant distinction between novel and familiar object evoked responses, and also a systematic change where novel patterns of activity are more dissimilar and familiar more similar (as indicated by lower dissimilarity, or the cooler shade in the bottom left "familiar" square in each RDM).

3.4 Discussion

In the current study, we explored whether object-based recognition memory signals in the MTL reflect the category or domain to which the object belongs, and whether any such category or domain specific recognition memory signals were present specifically in structures that represented that category or domain in response to stimuli presentation. Further, we asked whether or not recognition memory signals were coded in an overall change in signal (repetition suppression), or a pattern based change (pattern of increases and decreases). We tested the first question in two ways. First, we isolated the memory signal by contrasting novel and familiar objects, and tested whether the representational space of the memory signal for each object category was organized by category and domain. Second, we tested whether we could distinguish patterns of activity evoked by perceived novel or familiar objects within category, within domain, or across all categories. We found no evidence for category or domain specific memory signals. We did, however, find that it was possible to distinguish novel from familiar objects by patterns of activity evoked across all object categories, in all three medial temporal lobe structures. Furthermore, while we observed significant repetition suppression in all MTL structures examined, the distinction between patterns of activity for novel and familiar objects was not simply the result of a generic average repetition suppression effect. Instead, we found patterns of activity for novel and familiar objects that could be distinguished even after an average signal change was removed, and, interestingly, there is a systematic change in similarity, with novel objects being more dissimilar than familiar objects.

To our knowledge, our first analyses, in which we attempted to isolate the memory component of the signal, and explore the representational space of this signal, has not previously been done. Given this novelty, the fact that we did not find any category or domain organization using this approach is somewhat challenging to interpret. However, the findings from this analysis were in line with those from our second set of analyses which employed a more traditional approach. Specifically, in our second set of analyses we were unable to distinguish novel from familiar object-evoked responses within any of the categories or domains tested. This result was somewhat surprising, given category-specific memory signals in PrC and PhC have been reported for a number of the categories tested (faces, trees, buildings, and furniture); Martin et al., 2013, 2016). Importantly, while we did not find category or domain specific memory signals, we did find that novel and familiar patterns of activity could be distinguished - but only *across* object categories. This differing pattern of results could be due to a number of factors, or differences between experiments. One stark difference in design is that Martin et al. (2013, 2016) used a small number of object categories, with high numbers of within category exemplars, which presumably led to higher interference as evidenced by lower recognition memory performance; indeed d’ was 0.67 on average in the first study, and 0.59 on average in the second study. In contrast, in our task we employed exemplars from 12 different object categories, and, as a consequence, participants had relatively high recognition-memory performance (d’ =1.85). Another difference between studies is the analysis method. While Martin et al. showed that it was possible to classify perceived novel and familiar stimuli within but not between category, using a support vector machine, we used representational similarity analysis to test whether within category similarity was greater than between category similarity. However, generally speaking, these analyses typically converge, given that higher within versus between category similarity would allow for classification. A final, and potentially key difference, is that novel and familiar stimuli were necessarily different in Martin et al.’s studies which employed a study-test paradigm, whereas in the current experiment, we used a continuous recognition memory paradigm with identical stimuli.

In a study-test paradigm, because the novel stimuli are necessarily different from the repeated old stimuli, there is always the possibility that classification of novel and familiar stimuli is above chance, due to the classifier picking up on low-level differences between stimuli subsets, as opposed to the memory signal itself. A recent study by Huffman et al. (2017) provides some evidence that low-level confounds can potentially drive MVPA based findings that are thought to reflect mnemonic representations in the MTL. Specifically, patterns of activity in MTL evoked when subjects perform a standard associative memory task showed distinctions between contexts, items-in-context, and the temporal order of items-in-context. However, these seemingly high-level representations were also found in V1, and representations in both MTL and V1 showed a significant relationship to behavioral performance. Further, when low-level stimuli differences between items and contexts were highly controlled, these distinctions disappeared in MTL regions. One possibility is that the memory signal distinctions reported by Martin et al. were driven by similar low-level stimulus differences, an effect that might have disappeared when comparisons are made between identical new and old stimuli. However, stimuli employed in the Martin et al. tasks were controlled for a number of low-level features, and were counterbalanced across study and test, making this less likely. Additionally, in the studies by Martin et al., recognition memory performance was correlated with classification *exclusively* when classification included voxels that showed increases and decreases between novel and familiar stimuli, despite the fact that condition could be classified by using voxels that showed only increases or decreases (Martin et al., 2016). Overall, it is, thus, not clear that the differences across novel and familiar stimuli reported in these prior studies can be accounted for by low-level stimulus features. However, further experimentation is needed, such as testing whether the current pattern of results is present in V1. In particular, one could test the finding that the pattern of increases and decreases distinctly correlates with recognition memory performance.

An alternative explanation of the differing pattern results between studies is that Martin et al.'s task drives MTL regions to compute memory in a more categorical manner, while our task drives the MTL to compute memory in a more general manner. Or, rather, the different tasks may influence different aspects of the representational space, one that reflects the space evoked by seeing many diverse objects, and the other that reflects the aspect of space evoked by seeing many highly similar objects from a small number of categories.

In order to examine this possibility, it may be helpful to consider computational models of object processing in cortex, and how different models predict different results in terms of their category-specificity, as well as in terms of overall repetition suppression, or pattern based changes in coding. Representational-hierarchical computational models of PrC (Bussey & Saksida, 2002; Cowell et al., 2006) and complementary learning systems (CLS) models of medial temporal lobe cortex (MTLC) (Norman & O'Reilly, 2003), which share some commonalities in architecture and processing, can help us understand the differing memory signals in these two sets of results. In both models, objects are composed of conjunctive features so that when the network is initially presented with novel objects, these novel objects activate a large number of units (which can be conceptualized as neurons, or clusters of neurons) spread across the network (or MTL structure). As the same objects are repeated, through a process of Hebbian learning, the object representations are tuned, or a smaller number of local units are activated to a greater degree. This process is referred to as “sharpening” (Cowell et al., 2006). Through this process, the network becomes topologically organized - similar objects (those with overlapping features) are represented more similarly to each other in terms of shared units. In these models, the degree of sharpening is read out as a measure of stimulus familiarity - or is the recognition memory signal itself.

The link between recognition memory signals and stimuli are explicit in both of these models - the memory signal is a direct readout from the object representation. Because object representations are organized topologically (although to differing degrees in the CLS model and R-H models) this gives a precise explanation as to why memory signals can be categorical in nature (Cowell, 2006; Normal & O’Reilly, 2003). Indeed, if different object categories are represented distinctly in the network topology, it is understandable why training a classifier on activations of novel or familiar stimuli from one category (picking up on the pattern relevant to that specific topology or neighborhood) would not necessarily transfer to another category. Distinctions between novel and familiar stimuli can potentially be accounted for by the suggestion that novel stimuli activate a large number of distributed units, while objects, once repeated have a sharpened, and active a more distinct local topology (presumably these differences are what the classifier taps into). This latter explanation can potentially explain differences between novel and familiar stimuli within the same category *or* across all object categories. For example, the distinction between novel evoked responses that activate many units broadly spread throughout the network, and familiar evoked responses that activate a smaller number of units more strongly, may be reflected at both stimulus specific and stimulus general levels. If these interpretations of the models are correct, and there are both general and category-specific distinctions between novel and familiar stimuli, it remains an open question as to why Martin et al.’s tasks were sensitive to the category specific distinction, while our task was sensitive only to more general distinctions across category. It should be noted that in our initial analyses presented in Chapter 2, we do pick up on category and domain organization, which could be interpreted as picking up on the stimulus specific topology of the network. We no longer see this organization specifically when we evaluate memory signals.

In addition to shedding light on the category specificity or generality of recognition memory signals, these models also provide a mechanistic account of repetition suppression. Specifically, repetition suppression is a result of the decrease in the number of units responding when an item is repeated, so despite the fact that there is an increase in activity in some units, because so many units drop-out, there is an average decrease in activity in the network. This description of sharpening also fits the pattern based finding, described in Martin et al. (2016) in which voxels that show the largest changes, in terms of increases or decreases, correlate with better memory performance, given it is the pattern of change across the network that reflects sharpening. This finding and account of the pattern based changes also generally aligns with our findings that there are changes in patterns of activity beyond an average repetition suppression effect between novel and familiar object representations. In a sharpening model, there are increases and decreases in units that form a distinct pattern across the network during novel and familiar objects presentations in addition to the average signal decrease, mirroring our pattern of results where we see both an average signal change and a pattern based change beyond this. In other words, our results can potentially be thought of as tapping into these two aspects of the sharpening process. One note of caution when interpreting novel and familiar object evoked responses in the tasks discussed in light of these models, is that these models are designed to predict how a network organizes and responds when objects are FIRST learned and repeated (or are truly novel). In our paradigm, while the exact exemplars used are novel to participants, they already have familiarity in general with the object categories. Therefore, they presumably have some organization already present within the network for that object category based on previous experience.

3.4.1 Conclusion

In the current set of analyses, we did not find evidence of category or domain specificity in recognition memory signals in the MTL. This is somewhat surprising, given we see category and domain specificity in object stimuli evoked responses (Chapter 1), and previous work has shown category specific memory signals in MTL regions (Martin et al., 2013, 2017). One possibility is that we have decreased power in modeling our responses when split by novel and familiar (i.e., in chapter one we modeled all object evoked responses for each category regardless of whether it was novel or familiar, so our GLM estimates are based on more data points than in the current analyses). One straight-forward way to examine this is to test category and domain specificity within novel or familiar object trials only, or to do permutations on a smaller subsample of stimuli (the same size as currently used). Alternatively, our task, which involved many object categories, may have tapped into a general recognition memory signal in the network. Suggesting that at least in this task context, and perhaps others, recognition memory signals in the MTL may be computed (or picked up) as general signals across the network, somewhat distinct from the stimuli themselves.

In terms of how memory is coded, we replicated the finding that medial temporal lobe regions show robust decreases in average activity for repeated objects, previously reported at longer time lags (see meta-analysis: Kim, 2011). Moreover, beyond this repetition effect based on average activity, we showed that patterns evoked by repeated objects are distinct from patterns evoked by novel objects. These findings are generally in line with previous work showing that increases and decreases in voxels (beyond an average decrease or increase) best correlate with recognition memory (Martin et al., 2016). They are also in line with findings that neural changes with repetition in medial temporal areas go beyond a general scaling effect seen in more posterior VVS regions, and also can involve changes at the pattern level (Weiner et al., 2010).

Beyond finding a pattern-based change, we see a systematic change in how medial temporal lobe regions represent objects based on memory status, in which activity for familiar objects were more similar to each other overall than patterns of activity evoked by novel objects. This differs from previous work reporting changes that follow a “sharpening” pattern - where categories become *more distinct* from each other with repetition (Weiner et al., 2010). Therefore, it is curious that our findings show the opposite pattern - patterns evoked by objects from all categories become *more similar* with repetition. One possibility is that in our continuous recognition task with a large number of objects biases the MTL towards a more general distinction between novel and familiar object representations, perhaps because we are picking up a global pattern matching signal at retrieval that reflects global activity shared by all object representations in the network. Future work is needed to elucidate whether memory signals in the medial cortex are category specific, or more general, and whether evidence for either is based on task context. Finally, the relationship between how these either general or category specific memory signals are coded, whether it be changes in the average level of activity, in patterns of activity, or in both, may aid in gaining a deeper understanding of the neural mechanisms of object familiarity.

3.5 References

Aggleton, J. P., Brown, M. W., & Albasser, M. M. (2012). Contrasting brain activity patterns for item recognition memory and associative recognition memory: insights from immediate-early gene functional imaging. *Neuropsychologia*, *50*(13), 3141-3155.

Brown, M. W., & Xiang, J. Z. (1998). Recognition memory: neuronal substrates of the judgement of prior occurrence. *Progress in neurobiology*, *55*(2), 149-189.

Brown, M. W., & Aggleton, J. P. (2001). Recognition memory: what are the roles of the perirhinal cortex and hippocampus?. *Nature Reviews Neuroscience*, *2*(1), 51-61.

Bussey, T. J., & Saksida, L. M. (2005). Object memory and perception in the medial temporal lobe: an alternative approach. *Current opinion in neurobiology*, *15*(6), 730-737.

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

Bussey, T. J., Saksida, L. M., & Murray, E. A. (2002). Perirhinal cortex resolves feature ambiguity in complex visual discriminations. *European Journal of Neuroscience*, *15*(2), 365-374.

Cowell, R. A., Bussey, T. J., & Saksida, L. M. (2006). Why does brain damage impair memory? A connectionist model of object recognition memory in perirhinal cortex. *Journal of Neuroscience*, *26*(47), 12186-12197.

Cusack R, Vicente-Grabovetsky A, Mitchell DJ, Wild CJ, Auer T, Linke AC, Peelle JE (2015) Automatic analysis (aa): Efficient neuroimaging workflows and parallel processing using Matlab and XML. Frontiers in Neuroinformatics 8:90.http://dx.doi.org/10.3389/fninf.2014.00090.

Davis, T., Xue, G., Love, B. C., Preston, A. R., & Poldrack, R. A. (2014). Global neural pattern similarity as a common basis for categorization and recognition memory. *Journal of Neuroscience*, *34*(22), 7472-7484.

Duke, D., Martin, C. B., Bowles, B., McRae, K., & Köhler, S. (2017). Perirhinal cortex tracks degree of recent as well as cumulative lifetime experience with object concepts. *Cortex*, *89*, 61-70.

Eichenbaum, H., Yonelinas, A. P., & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annu. Rev. Neurosci.*, *30*, 123-152.

Fahy, F. L., Riches, I. P., & Brown, M. W. (1993). Neuronal activity related to visual recognition memory: long-term memory and the encoding of recency and familiarity information in the primate anterior and medial inferior temporal and rhinal cortex. *Experimental Brain Research*, *96*(3), 457-472.

Frankó, E., Insausti, A. M., Artacho-Pérula, E., Insausti, R., & Chavoix, C. (2014). Identification of the human medial temporal lobe regions on magnetic resonance images. *Human brain mapping*, *35*(1), 248-256.

Gonsalves, B. D., Kahn, I., Curran, T., Norman, K. A., & Wagner, A. D. (2005). Memory strength and repetition suppression: multimodal imaging of medial temporal cortical contributions to recognition. *Neuron*, *47*(5), 751-761.

Graham, K. S., Barense, M. D., & Lee, A. C. (2010). Going beyond LTM in the MTL: a synthesis of neuropsychological and neuroimaging findings on the role of the medial temporal lobe in memory and perception. *Neuropsychologia*, *48*(4), 831-853.

Grill-Spector, K., & Weiner, K. S. (2014). The functional architecture of the ventral temporal cortex and its role in categorization. *Nature Reviews Neuroscience, 15*(8), 536-548.

Heusser, A. C., Awipi, T., & Davachi, L. (2013). The ups and downs of repetition: Modulation of the perirhinal cortex by conceptual repetition predicts priming and long-term memory. *Neuropsychologia*, *51*(12), 2333-2343.

Huffman, D. J., & Stark, C. E. (2014). Multivariate pattern analysis of the human medial temporal lobe revealed representationally categorical cortex and representationally agnostic hippocampus. *Hippocampus*, *24*(11), 1394-1403.

Konkle, T., & Caramazza, A. (2013). Tripartite organization of the ventral stream by animacy and object size. *The Journal of Neuroscience, 33*(25), 10235-10242.

Kim, H. (2013). Differential neural activity in the recognition of old versus new events: An Activation Likelihood Estimation Meta-Analysis. *Human Brain Mapping*, *34*(4), 814-836.

Kivisaari, S. L., Tyler, L. K., Monsch, A. U., & Taylor, K. I. (2012). Medial perirhinal cortex disambiguates confusable objects. *Brain*, *135*(12), 3757-3769.

Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, *2*, 4.

Kriegeskorte, N., Mur, M., Ruff, D. A., Kiani, R., Bodurka, J., Esteky, H., ... & Bandettini, P. A. (2008). Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron*, *60*(6), 1126-1141.

LaRocque, K. F., Smith, M. E., Carr, V. A., Witthoft, N., Grill-Spector, K., & Wagner, A. D. (2013). Global similarity and pattern separation in the human medial temporal lobe predict subsequent memory. *Journal of Neuroscience*, *33*(13), 5466-5474.

Lee, A. C., Scahill, V. L., & Graham, K. S. (2007). Activating the medial temporal lobe during oddity judgment for faces and scenes. *Cerebral Cortex*, *18*(3), 683-696.

Li, L., Miller, E. K., & Desimone, R. (1993). The representation of stimulus familiarity in anterior inferior temporal cortex. *Journal of neurophysiology*, *69*(6), 1918-1929.

Liang, J. C., Wagner, A. D., & Preston, A. R. (2013). Content representation in the human medial temporal lobe. *Cerebral Cortex*, *23*(1), 80-96.

Linke, A. C., Vicente-Grabovetsky, A., & Cusack, R. (2011). Stimulus-specific suppression preserves information in auditory short-term memory. *Proceedings of the National Academy of Sciences*, *108*(31), 12961-12966.

Litman, L., Awipi, T., & Davachi, L. (2009). Category-specificity in the human medial temporal lobe cortex. *Hippocampus*, *19*(3), 308-319.

Martin, C. B., McLean, D. A., O'Neil, E. B., & Köhler, S. (2013). Distinct familiarity-based response patterns for faces and buildings in perirhinal and parahippocampal cortex. *The Journal of Neuroscience*, *33*(26), 10915-10923.

Martin, C. B., Cowell, R. A., Gribble, P. L., Wright, J., & Köhler, S. (2015). Distributed category-specific recognition-memory signals in human perirhinal cortex. *Hippocampus*.

Martin CB, Sullivan JS, Wright J, Köhler S. (2017). Recognition-memory signals for objects from different categories are graded across perirhinal and parahippocampal cortex. (under review).

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

Moss, H. E., Rodd, J. M., Stamatakis, E. A., Bright, P., & Tyler, L. K. (2004). Anteromedial temporal cortex supports fine-grained differentiation among objects. *Cerebral cortex*, *15*(5), 616-627.

Pruessner, J. C., Li, L. M., Serles, W., Pruessner, M., Collins, D. L., Kabani, N., ... & Evans, A. C. (2000). Volumetry of hippocampus and amygdala with high-resolution MRI and three-dimensional analysis software: minimizing the discrepancies between laboratories. *Cerebral cortex*, *10*(4), 433-442.

Norman, K. A., & O'reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. *Psychological review*, *110*(4), 611.

Pruessner, J. C., Köhler, S., Crane, J., Pruessner, M., Lord, C., Byrne, A., ... & Evans, A. C. (2002). Volumetry of temporopolar, perirhinal, entorhinal and parahippocampal cortex from high-resolution MR images: considering the variability of the collateral sulcus. *Cerebral Cortex*, *12*(12), 1342-1353.

Ringo, J. L. (1996). Stimulus specific adaptation in inferior temporal and medial temporal cortex of the monkey. *Behavioural brain research*, *76*(1), 191-197.

Segaert, K., Weber, K., de Lange, F. P., Petersson, K. M., & Hagoort, P. (2013). The suppression of repetition enhancement: a review of fMRI studies. *Neuropsychologia*, *51*(1), 59-66.

Walther, A., Nili, H., Ejaz, N., Alink, A., Kriegeskorte, N., & Diedrichsen, J. (2016). Reliability of dissimilarity measures for multi-voxel pattern analysis. *Neuroimage*, *137*, 188-200.

Wang, W. C., Ranganath, C., & Yonelinas, A. P. (2014). Activity reductions in perirhinal cortex predict conceptual priming and familiarity-based recognition. *Neuropsychologia*, *52*, 19-26.

Xiang, J. Z., & Brown, M. W. (1998). Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology*, *37*(4), 657-676.

Yassa, M. A., & Stark, C. E. (2008). Multiple signals of recognition memory in the medial temporal lobe. *Hippocampus*, *18*(9), 945-954.

Ye, Z., Zhu, B., Zhuang, L., Lu, Z., Chen, C., & Xue, G. (2016). Neural global pattern similarity underlies true and false memories. *Journal of Neuroscience*, *36*(25), 6792-6802.

Chapter 4

4 Abnormal semantic knowledge in a case of developmental amnesia

4.1 Introduction

Developmental amnesia can result from early extended hippocampal system pathology, often due to perinatal hypoxia, and is characterized by a pronounced deficit in episodic recollection starting early in life. In addition to a lack of episodic recollection, individuals with developmental amnesia present with severe difficulties navigating in familiar environments, remembering where items are placed, and orienting themselves by date and time (Vargha-Khadem et al., 1997). In clinical neuropsychological tests of long-term memory, individuals with developmental amnesia score far below the general population. However, despite having severe episodic memory impairment, these individuals typically have normal IQs and progress successfully through mainstream schooling (Vargha-Khadem et al., 1997). Indeed, individuals with developmental amnesia typically show average academic performance, as well as normal performance on reading and spelling tests, standardized vocabulary tests, and general information questions, such as “What is the capital of France?” These observations have been taken as evidence that semantic knowledge is preserved in developmental amnesia (Vargha-Khadem et al., 1997; Gadian et al., 2000; Bindschaedler et al., 2011; Brizzolara, 2003). This view has received further support from experimental studies showing that new semantic knowledge can be acquired in developmental amnesia (Guillery-Girard et al., 2004; Bindschaedler et al., 2011; Martins et al., 2007; but see Gardiner et al., 2008).

Episodic and semantic memory are often considered to be separate, dissociable memory systems (Schacter and Tulving, 1994), and core retrieval processes from each system are thought to be independent. Despite being dissociable, there is a dependence between the two systems with respect to encoding of information; however, there are two *distinct* models of the direction of this dependence. According to the serial-parallel-independent (SPI) model (Tulving & Markowitsch, 1998; Tulving, 2002), episodic memory is a phylogenetically and ontogenetically later-developing system that grew out of semantic memory. As such, encoding episodic information critically depends on semantic memory, whereas semantic memory encoding is independent of episodic memory. By contrast, models of memory transformation or consolidation (Winocur et al., 2010; McClelland et al., 1995) suggest that semantic information is encoded initially as hippocampal-dependent episodic memory and then, either through repetition or gist extraction, becomes semantic. Evidence of intact semantic memory in developmental amnesia is considered one of the strongest sources of support for the SPI claim that semantic memory can be acquired independently of episodic memory (Vargha-Khadem, 1997; Tulving, 2002).

The extent to which semantic and episodic memory are dissociable at the neural level has been widely debated, with some researchers arguing that both semantic and episodic memory are dependent on the integrity of the medial temporal lobes, including the hippocampus (Squire, 1987; Shimamura & Squire, 1987; Gabrieli et al., 1988; Zola-Morgan, Cohen, & Squire, 1983; Squire & Zola, 1998). In contrast, others have argued that the medial temporal lobes are crucial for episodic but not semantic memory (Kinsbourne & Wood, 1975; Parkin, 1982). This latter view was supported largely by the finding that individuals with amnesia due to hippocampal damage seemed to show primarily deficits in episodic memory, with semantic memory being relatively preserved. The finding that individuals with developmental amnesia due to focal hippocampal damage had intact semantic memory added strong support to the notion that episodic but not semantic memory is dependent on the hippocampus.

Since the publication of Vargha-Khadem (1997), the prevailing view has been that semantic memory is normal in developmental amnesia. However, some interesting variations have been noted. For example, a few studies report that patients can acquire new semantic information, but that they require a greater number of repetitions than do control participants (Gardiner et al., 2008; Guillery-Girard et al., 2004). Furthermore, one patient did not acquire facts about the world at the same rate as normal control participants (Bindschaedler et al., 2011). Another individual with developmental amnesia showed below average performance on the “information” I.Q. subtest (although normal performance on a questionnaire about world facts), as well as poor naming and comprehension abilities (Vicari et al., 2007). These findings hint at possible subtle abnormalities in the development of semantic memory against a background of severe episodic memory impairment. However, it is critical to understand the nature of these abnormalities for a clearer view of the relative contribution of episodic memory and hippocampal function to the early formation of semantics.

An important component of semantic memory that has yet to be explored in developmental amnesic cases is the structure of their conceptual knowledge. This can be explored through knowledge of concrete concepts such as “chair” and “dog”, and we focus on these types of concepts in this article. One of the dominant theoretical frameworks for understanding concrete concepts is feature-based models in which concepts are represented in terms of semantic features (Martin, 2007; Tyler et al., 2013). In these models, theoretically and empirically important distinctions have been made among the types of features of which concepts may be composed (Cree & McRae, 2003; Wu & Barsalou, 2009). Consider “hammer”. Learning the concept of hammer can be thought of in terms of learning its features, such as how it looks (elongated, has a handle), how it is used (grip the handle, swing back and forward), what it is used for (pounding nails into the wall), and where it is commonly found (the garage). Some aspects of people’s knowledge of objects such as hammers are intrinsic to, or physically part of, the object itself, whereas other aspects can be considered to be part of the contexts in which an object is encountered (Barr & Caplan, 1987). Contextually based information may include how or where a hammer is used, and can be considered as extrinsic to the object. For concrete living (cow) and nonliving things (hammer), important aspects of extrinsic information involve relations between the physical object itself and other types of objects and locations, as well as the ways in which people interact with the object.

During development, individuals acquire both intrinsic and extrinsic feature knowledge as part of learning concrete concepts. In terms of neural mechanisms, the hippocampus may play an important role in binding extrinsic features to the intrinsic features that compose the object. A large body of evidence suggests that *relational binding*, or the forming of associations among an object and its context, is a core function of the hippocampus (Olsen et al., 2012; Eichenbaum & Cohen, 1993). Anatomically, the hippocampus is well suited to the task because it receives object information from the visual ventral stream, as well as contextual information (e.g., found in a garage) from a number of unimodal and multimodal cortical areas (Lavenex & Amaral, 2000). Although typically this type of learning is considered with reference to individual episodes, some evidence suggests that acquisition of new conceptual knowledge may also rely on hippocampally mediated binding, particularly across multiple episodes that share regularities in object-context relationships (Kumaran et al., 2009). To the extent that developmental amnesia is the result of damage to the hippocampus or the extended hippocampal system in most, if not all, cases, one might expect extrinsic aspects of conceptual knowledge to develop abnormally.

If extrinsic aspects of concept knowledge are impaired in developmental amnesia, how might this affect the structure of semantic memory? There exists a correspondence between Barr and Caplan’s (1987) distinction between intrinsic and extrinsic features and the distinction that has been made between sensory and functional features (or sensory and non-sensory features). Sensory information corresponds to intrinsic knowledge because people’s (somewhat abstract) knowledge of how something looks, smells, and sounds is intrinsic to the entity or object itself. In contrast, functional information reflects extrinsic information because how an object is used and what it is used for are based on relationships between the object and something else, such as actions performed on the object by some agent. In some cases, the term functional information has been used to refer to all knowledge about a concept that is not sensory (Warrington & Shallice, 1984). In this sense, in addition to information about how an object is used and what it is used for, functional (or actually “nonsensory”) information includes other extrinsic information, such as where an entity or object typically is located, when an object typically is used (e.g., during the winter), or what an object typically is used with. Note that researchers differ with regard to how various types of information are labelled or classified, and detailed classification schemes have been proposed and used (Cree & McRae, 2003; Wu & Barsalou, 2009). For our present purposes, what is centrally important is that sensory information is intrinsic to an object itself, whereas other types of knowledge about concrete concepts are extrinsic (Barr & Caplan, 1987).

Interestingly, a recent study showed that patient D.A., an individual with adult-onset amnesia associated with hippocampal damage, was impaired in learning new functional knowledge in an experimental setting. Specifically, although D.A. learned movements related to new, unfamiliar tools as quickly as control participants, he was subsequently unable to recall the functional significance of those objects (what they were used for) as well as controls (Roy & Park, 2010). An intriguing possibility, therefore, is that for concrete concepts, the learning of extrinsic features outside the laboratory also depends on hippocampal integrity, and may be impaired in developmental amnesia.

Our goal was to test whether some aspects of semantic knowledge depend on the same relational binding mechanisms in the hippocampus that subserve episodic memory. Specifically, we explored whether extrinsic feature knowledge for concrete concepts is impaired in developmental amnesia. We investigated a previously described individual with developmental amnesia, HC, with well-documented impairments in episodic memory and a well-characterized lesion of the extended hippocampal system (Olsen et al., 2013). We tested HC on a semantic feature production task and a typicality rating task, both of which are sensitive to knowledge that underlies individual concepts. Although these tasks have been used frequently in the concepts literature on healthy adults, to our knowledge they have not been employed in research on semantic memory in developmental amnesia. In order to evaluate the anatomical specificity of any abnormalities, we also compared HC with an individual with anterior temporal lobe damage that spares the hippocampus, patient NB (Bowles et al., 2007; Bowles et al., 2016). We predicted that any abnormalities HC showed on these tasks would result from her hippocampal lesion, and therefore we predicted that NB would not show any of the same abnormalities in our experimental tasks.

In Experiment 1, we used a semantic feature production task to test whether HC would produce fewer extrinsic features than would control participants, while producing similar numbers of intrinsic features. We classified the features generated by participants using Cree and McRae’s (2003) feature-type taxonomy. A feature type was counted as an extrinsic feature if it described how the object is associated to other entities, objects, locations, and so forth, rather than being intrinsic to the object itself. Extrinsic features included functional, encyclopedic, and taxonomic features, whereas intrinsic features included smell, taste, sound, colour, visual form and surface, and visual motion features. Note that tasks that require explicit feature generation may be sensitive to both abnormal acquisition of knowledge, and retrieval-related impairments. Because it has been shown that constructing relational conceptual representations ‘online’ is impaired in adult-onset amnesia (Rosenbaum et al., 2009; Waidergoren et al., 2012), we therefore included a second experiment that did not require explicit generation of feature names.

In Experiment 2, we used a typicality rating task to test whether differences in knowledge of extrinsic versus intrinsic features results in HC producing abnormal ratings for nonliving things, but not for living things. Typicality ratings require participants to judge the “goodness” of a concept as a member of a particular category. For example, on a 1 to 9 scale in which 9 corresponds to extremely typical of a category, a robin might be given a rating of 9 for the category of bird, whereas a penguin might be given a 2. In many theories, rating typicality for concrete concepts involves comparing feature-based knowledge of the lower-level concept (such as *hammer*) to either the representation of the higher-level concept (such as *tool*; Rosch & Mervis, 1975), or to other concepts within the same category (other tools; Medin & Schaffer, 1978). Intrinsic and extrinsic features are differentially weighted when people rate typicality for categories that are drawn from living and nonliving domains, respectively. Barr and Caplan (1987) showed that intrinsic features are primarily important for determining category membership for living things, whereas extrinsic (relational) features are primarily important for determining category membership for nonliving things (see also Keil, 1989). Similarly, it has been shown that function (which is extrinsic and relational) is central for adults when categorizing nonliving things, whereas appearance and behavior (both of which are intrinsic) are central for categorizing living things (Barton & Komatsu, 1989; Hampton et al., 2009). There also is developmental research that speaks to these issues. For example, Kalénine and Bonthoux (2008) conducted a triad task (which of 2 pictured objects goes with the base object?) with children 5 and 7 years of age, as well as with adults. All three groups had shorter latencies when using contextual/functional relationships for manipulable objects, and when using perceptually-based relations for living things. A number of studies also have shown that functional similarities strongly influence categorization, induction, and name generalization for nonliving things from at least age 2 onward (Casler & Kelemen, 2007; Diesendruck, Markson, & Bloom, 2003; Kemler Nelson, Chan Egan, & Holt, 2004; Truxaw, Krasnow, Woods, & German, 2006). In contrast, children depend on perceptual similarity to categorize living things (Sloutsky, Kloos, & Fisher, 2007; Sloutsky & Spino, 2004). Thus, from a very early age, children learn and use intrinsic and extrinsic knowledge of object concepts, and they are sensitive to the correspondence between these types of information and the living/nonliving distinction. Rating typicality is similar to categorization in that it can be considered a process of rating the ease with which something can be characterized as belonging to a category. Therefore, we used a typicality rating task with living and nonliving thing categories because it has been shown that these ratings differentially depend on knowledge of intrinsic and extrinsic features. Furthermore, because it is not a binary judgment, it provides a sensitive measure. We predicted that HC's typicality ratings would be abnormal for nonliving things because these ratings depend primarily on extrinsic features. In contrast, HC's typicality ratings for living things should be normal because they are based primarily on intrinsic features.

With respect to NB, who has shown deficits in familiarity judgments, an advantage of using typicality ratings is that concept familiarity minimally influences typicality ratings. Of course, familiarity can influence typicality ratings in the sense that if a person has no knowledge of the concept, she will rate it as being atypical (Malt & Smith, 1982). However, the one study that directly investigated the relationship between rated concept familiarity and typicality found no relation between the two (Barsalou, 2003). This makes sense in that familiarity is not linked to any particular category, whereas typicality is. For example, *lamp* is an atypical member of *furniture*, but *lamp* itself is a highly familiar concept. Furthermore, *snake*, regardless of its familiarity, is a typical *reptile*, but an atypical *animal*. In general then, typicality ratings for concepts of common superordinate categories like those used in Experiment 2 are strongly influenced by aspects of conceptual structure, but are influenced by familiarity to only a minor extent (Barsalou, 1985; Malt & Smith, 1982; Rosch, Simpson, & Miller, 1976).

4.2 Participants

4.2.1 HC

HC is a well-documented developmental amnesic person who has participated in several previous studies (Adlam et al., 2005, 2009; Kwan et al., 2010; Maguire et al., 2011; Olsen et al., 2013; Rabin et al., 2012; Rosenbaum et al., 2011, 2015; Vargha-Khadem et al., 2003). She is a right-handed female who was 22 years old at the time of testing. HC finished high school, one year of technical college, and an additional year in a post-secondary vocational program (total education = 14 years). It initially was presumed that HC suffered hypoxia perinatally as a result of premature birth (gestational age = 32 weeks; Vargha-Khadem et al., 2003). Her lesion has been characterized extensively in past studies, revealing bilateral volume reduction in the hippocampus (see Figure 4.1). High-resolution MR-based volumetric analyses of HC’s medial temporal-lobe structures revealed significantly reduced hippocampal volumes in anterior sections bilaterally and additional volume reductions in the right posterior hippocampus, as compared to controls matched in sex, age, and education (Olsen et al., 2013). Overall, hippocampal volume was reduced by approximately 30%. Critically, perirhinal (PRC), entorhinal (ERC), and parahippocampal (PHC) cortices were fully preserved and showed no volume reduction. In a follow-up study, it was found that HC has abnormal morphology of the fornix, absent mammillary bodies, and inverted hippocampal structure (Rosenbaum et al., 2014), suggesting that a prenatal etiology may account for her developmental amnesia.

HC shows clear signs of amnesia with pronounced impairments on clinical neuropsychological tests of long-term memory, experimental tasks of recognition memory, and autobiographical memory tasks (see Table 4.1 for a comprehensive neuropsychological profile). To illustrate, on the California Verbal Learning Test, HC's impaired performance is reflected in a *z*-score of -4 for short delay free recall, -3 for long-delay free recall, and -2 for long-delay recognition memory. Testing of autobiographical memory using the Autobiographical Interview (AI) and the Galton-Crovitz paradigms revealed that HC’s autobiographical recollections lack episodic details across her life-span (for further detail, see Kwan et al., 2010; Rabin et al., 2012; Rosenbaum et al., 2011). Despite her long-term memory deficits, HC appears to have normal intelligence and semantic memory. Specifically, her WASI Full Scale IQ is in the 66th percentile, and semantic fluency, as tested using animals, is above the 90th percentile (see Rosenbaum et al., 2011, for a full neuropsychological profile). Research with HC as a participant was part of a larger program (directed by R.S. Rosenbaum) that was approved by the Research Ethics Boards of York University and Baycrest in Toronto.

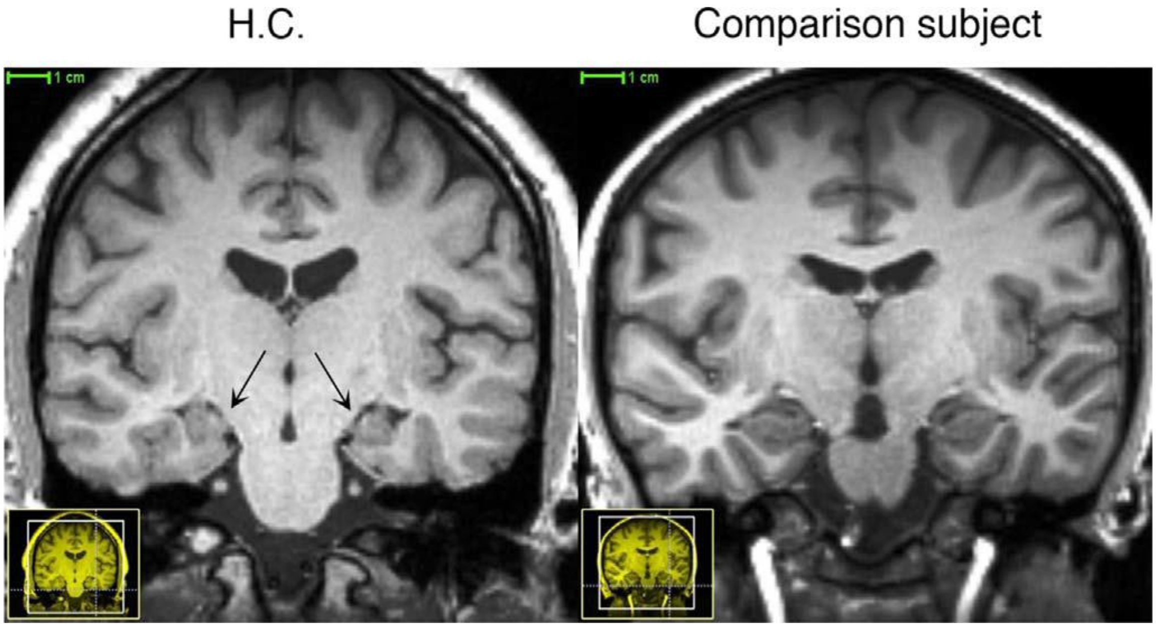

Figure 4.1. Coronal slice of T1-weighted MR image showing bilateral hippocampal volume reduction in HC and a representative control participant (Olsen et al., 2013).

Table 4.1: Neuropsychological Profile of Patient HC

Test	*Normed Score*
General Intellectual Function	
WASI	
Verbal IQ (percentile)	61st
Performance IQ (percentile)	66th
Full Scale IQ (percentile)	66th
AM-NART (standard score)	101.28 (estimated FSIQ)
Semantic Knowledge and Academic Attainment	
WASI Vocabulary (T-score)	55th
WAIS-III Information (scaled score)	12
WAIS-III Arithmetic (scaled score)	8
Language Production	
Boston Naming Test (z-score)	0.75
Semantic Fluency (animals)(percentile)	>90th

Anterograde Memory	
WMS-III	
Logical Memory I – immediate recall (scaled score)	4
Logical Memory II – delayed recall (scaled score)	1
California Verbal Learning Test-II	38
Total trials 1-5 (T-score)	
Short delay free recall (z-score)	-4
Long delay free recall (z-score)	-3
Recognition (z-score)	-2
Rey Osterreith complex figure	
Immediate recall (T-score)	<20
Delayed recall (T-score)	<20
Delayed recognition – total correct (T-score)	22
Processing Speed	
WAIS-III Digit Symbol (scaled score)	13
WAIS-III Symbol Search (scaled score)	14
Visuospatial Function	
WASI Block Design (T-score)	5

Rey-Osterrieth Complex Figure – Copy (percentile)	>16th
Judgement of Line Orientation (percentile)	56th
Benton Facial Recognition (percentile)	33-59th
Attention and Executive Function	
Stroop	
Word full (seconds) (z-score)	3.65
Color full (seconds) (z-score)	-0.03
Interference (seconds) (z-score)	-0.57
Word errors (z-score)	0
Colour errors (z-score)	-0.5
Interference errors (z-score)	-0.13
Word self-corrections (z-score)	-0.5
Colour self-corrections (z-score)	-0.71
Interference self-corrections (z-score)	1.44
Trail Making Test	
Part A (sec) (z-score)	0.69
Part B (sec) (z-score)	-0.23
Phonemic Fluency (FAS) (percentile)	70-80th

WASI Similarities (T-score)	50
WASI Matrix Reasoning (T-score)	55
Wisconsin Card Sorting Task – categories (T-score)	57
Mood	
CES-D	no indication of depression
PANAS	
Positive Affect (z-score)	1.06
Negative Affect (z-score)	0.21

4.2.2 NB

Like HC, NB is a well-documented individual who has taken part in a number of experiments which have carefully documented her memory deficits and lesion characteristics (Bowles et al., 2007; 2010; 2011; 2016, & Martin et al., 2011). NB was 26 at the time of testing, is right-handed, a native English speaker, and university educated. She underwent surgery to remove a mass in her left amygdala that was causing intractable epilepsy. Her surgical resection involved the most anterior portion of the medial and lateral left temporal lobe. Volumetric follow-up analyses showed that compared to the healthy right temporal lobe, the surgery resulted in removing large portions of her left amygdala, perirhinal cortex, and entorhinal cortex, sparing the hippocampus (Bowles et al., 2007). Additionally, evidence from an fMRI study suggests that her left hippocampus is functionally intact (Bowles et al., 2011). A post-surgical neuropsychological examination showed that NB has normal cognitive function in all domains, except for a low semantic fluency score (21st percentile) (see Supplementary Table 1 in Bowles et al., 2007).

4.2.3 Control participants

For HC, ten healthy control participants (Mean age = 22.1 years, *SD*= 1.9 years; 9 females), predominantly attending college at the time of testing, took part in the study. To minimize potential effects of achievement differences in tests of semantic memory, they were matched closely in education to HC in both years (Mean = 14.2 years, *SD* = 0.93 years) and level of education. It is important to note that while in many countries the terms college and university are used interchangeably, in Canada, “college” refers to a higher education institute in applied arts and technology that is more hands-on and career-oriented than a traditional university. Colleges usually offer either a 1 year certificate program, or a 2-3 year bachelor program. In the current study, predominantly college rather than university participants took part; 5 went to college only, 2 attended both college-level and university-level courses, and 3 attended university-level only. Controls were screened in a brief interview to rule out the presence of any history of neurological disorders. Control participant testing was approved by the Research Ethics Boards of York University and Baycrest, as well as the Fanshawe College Human Ethics Board. The same control participants were tested in Experiments 1 and 2.

Control participants for patient NB were 10 university-educated females (Mean age = 23.3, SD = 1.8). These control participants took part in a battery of tests as part of a separate study on NB, approved by the Health Sciences Research Ethics Board and the internal Ethics Board of the Department of Psychology at Western University.

4.3 Experiment 1: Feature production

4.3.1 Materials

The stimuli consisted of 40 concepts chosen from McRae, Cree, Seidenberg, and McNorgan (2005). There were 16 living things such as *pig* and *potato*, and 24 nonliving things such as *shoes* and *blender*. This number of concepts is approximately the upper limit that could reasonably be used for this intensive task, at least in a single session. By comparison, most semantic feature production studies have used approximately 20-25 concepts per participant.

4.3.2 Procedure

For each concept, participants were presented with a piece of paper with a concept name at the top, and were given 90 seconds to write down as many features as possible. The instructions were identical to those employed by McRae et al. (2005; see their Appendix B), although that study did not incorporate a time limit. Features were defined to participants as properties of the concept to which the word refers. They were further instructed that features could include: physical properties, such as internal and external parts; functional properties, such as what it is used for, as well as where, when, and by whom it is used; things that the concept is related to, such as the category that it belongs in; and other facts, such as how it behaves, or where it comes from. Participants were able to follow these instructions. For example, for squid, one control participant produced the following: <slimy>, < swims>, <eats fish>, <lives in the ocean>, <tentacles>, <floats in the water>, and <smaller than octopus>. The instructions remained in plain view throughout the task, and reminders of the kinds of features that could be provided were given periodically to avoid the potential impact of memory impairments for the instructions on performance. The experiment took approximately one hour.

4.3.3 Results

The listed features were coded using McRae et al.'s (2005) procedures. The features were coded by three independent judges who were blind to whether the features had been produced by HC, NB, or controls. The features were coded initially into knowledge types using Cree and McRae's (2003) brain region classification scheme. There are nine knowledge types in this classification scheme. Visual form and surface features denote parts of objects and entities (<has legs>, as well as shape (<is round>), and size (<is large>). Colour features denote colours of objects and entities (<is green>). Sound features describe the sounds that they make (<oinks>, <is loud>). Tactile features describe how an object or entity feels (<is soft>). There are also taste features (<tastes salty>). Visual motion features denote how an entity or object moves on its own (<swims>). Functional features denote how something is used, what it is used for, or who typically uses it (<used for repairs>). Encyclopedic features are somewhat less homogeneous than the other classes, although the majority of them convey information about location or time, such as <found on farms> or <worn in winter>. Finally, taxonomic features primarily denote the category (or categories) to which an object or entity belongs (<is an animal>), although some category co-ordinates and category exemplars were listed (<potbelly> for *pig*).

Our primary goal was to test whether HC demonstrates abnormalities in producing features that are extrinsic in nature. Because it often has been argued that the hippocampus encodes relations, or associations among types of information (Konkle & Cohen, 2012; Mayes et al., 2007) and extrinsic features are contextual or associative in nature, we hypothesized that HC would produce fewer extrinsic features for concrete concepts. Extrinsic features included functional, encyclopedic, and taxonomic features, whereas intrinsic features included smell, taste, sound, colour, visual form and surface, and visual motion features. The primary determinant of whether a feature type was counted as extrinsic was whether it described how the object is associated to other entities, objects, locations, and so forth, rather than being intrinsic to the object itself. Thus, learning these types of knowledge arguably relies on a relational binding mechanism.

We first examined the total number of features produced by HC, NB, and controls, irrespective of feature type, across the 40 concepts. HC produced 8.3 features per concept, which was non significantly fewer than controls (*M* = 11.1, *SD* = 3.4; *range* = 7.2 - 18.5), *z* = -0.82, *p* > .4. NB did not differ from controls in terms of overall number of features produced (*NB Mean* = 10.1, *control Mean* =10.4, SD = 1.6, *range* = 8.6 to 12.2 *z* = 0.19, *p* > .8).

To address the primary question of whether HC produces fewer extrinsic features than do controls, we averaged the number of features across all 40 concepts for HC and controls for each type of feature. HC produced on average 3.6 intrinsic features per concept (*SD* = 1.7, *range* = 0 - 8), which was almost identical to the 3.7 intrinsic features per concept (*SD* = 0.9, *range* = 2.1 - 6.1) produced on average by control participants. In terms of extrinsic features, however, HC produced 4.7 per concept (*SD* = 1.8, *range* = 2 - 10), whereas controls produced 6.9 per concept (*SD* = 1.2, *range* = 5 - 10). Numerically, these data are in line with our predictions that HC has selective difficulties with extrinsic features. We first analyzed these data by using items (the 40 concepts) as the replication factor. Using items as the replication factor enabled conducting a mixed analysis of variance using the number of features produced as the dependent variable. The independent variables were feature type (intrinsic vs. extrinsic), which was a within-items factor, and group (HC vs. controls), which was a between-items factor. Feature type and group interacted, $F(1,78) = 16.59$, $p < .0001$. Simple main effects analyses (using a pooled error term) showed that HC produced significantly fewer extrinsic features than did controls, $F(1,153) = 44.54$, $p < .0001$. The 0.1 features per concept difference for intrinsic features was nonsignificant, $F < 1$.

Next, we analyzed the feature data in a manner that tests the difference between HC and controls for intrinsic and extrinsic features by using difference scores. This permitted an analysis based on participants, rather than items. First, the average number of intrinsic and extrinsic features for all 40 concepts was calculated for HC and for each control. A difference score was calculated by subtracting the average number of features of each type (intrinsic and extrinsic) produced by HC from each control participant's average. For extrinsic features, HC produced 2.2 fewer features on average than did controls, which was significantly different from 0 using a single-sample *t*-test, $t(9) = 2.89$, $p < .02$. In contrast, for intrinsic features, HC produced only 0.6 fewer features on average than did controls, $t(9) = 1.35$, $p > .2$.

In contrast, NB showed no differences from her controls in feature production for either intrinsic or extrinsic features. NB produced on average 3.6 intrinsic features per concept ($SD = 1.8$), which was identical to the 3.6 intrinsic features per concept ($SD = 1.5$) produced on average by control participants. Importantly, unlike HC, NB produced 6.5 extrinsic features per concept ($SD = 1.9$), which is almost identical to the 6.7 extrinsic features per concept produced by her controls ($SD = 1.2$). There was no interaction between feature type and group $F(1,78) = 0.06$, $p > .8$. Additionally, there were no significant differences in the number of extrinsic or intrinsic features produced, when using NB’s difference score, NB produced 0.2 fewer extrinsic features ($t(5) = 0.37$, $p > .7$), and 0.1 fewer intrinsic features ($t(5) = 0.33$, $p > .7$).

Experiment 1 demonstrates that HC produced significantly fewer extrinsic features relative to controls. In contrast, HC produced virtually the same number of intrinsic features. Importantly, an individual with medial temporal damage sparing the hippocampus, NB, produced the same number of intrinsic and extrinsic features as controls. This supports our hypothesis that HC has specific abnormalities in her production of extrinsic knowledge of object concepts. That is, she lacks a rich store of extrinsic feature information, presumably because acquiring extrinsic knowledge relies on relational binding by the hippocampus. Note that this type of impairment in semantic knowledge would not be apparent in standard clinical tests of semantic memory. Consider, for example, that HC scored at the 90th percentile in category fluency when asked to produce names of animals. People’s knowledge of animals tends to be dominated by intrinsic knowledge, such as what an animal looks like, what it sounds like, or how it behaves on its own. On the other hand, people’s knowledge of many nonliving things such as types of tools or furniture, is dependent to a large degree on knowledge of how humans interact with these objects, or where they tend to be located. Because Experiment 1 provided a relatively direct measure of these types of knowledge, HC’s hippocampal-based impairment regarding extrinsic relational knowledge was evident.

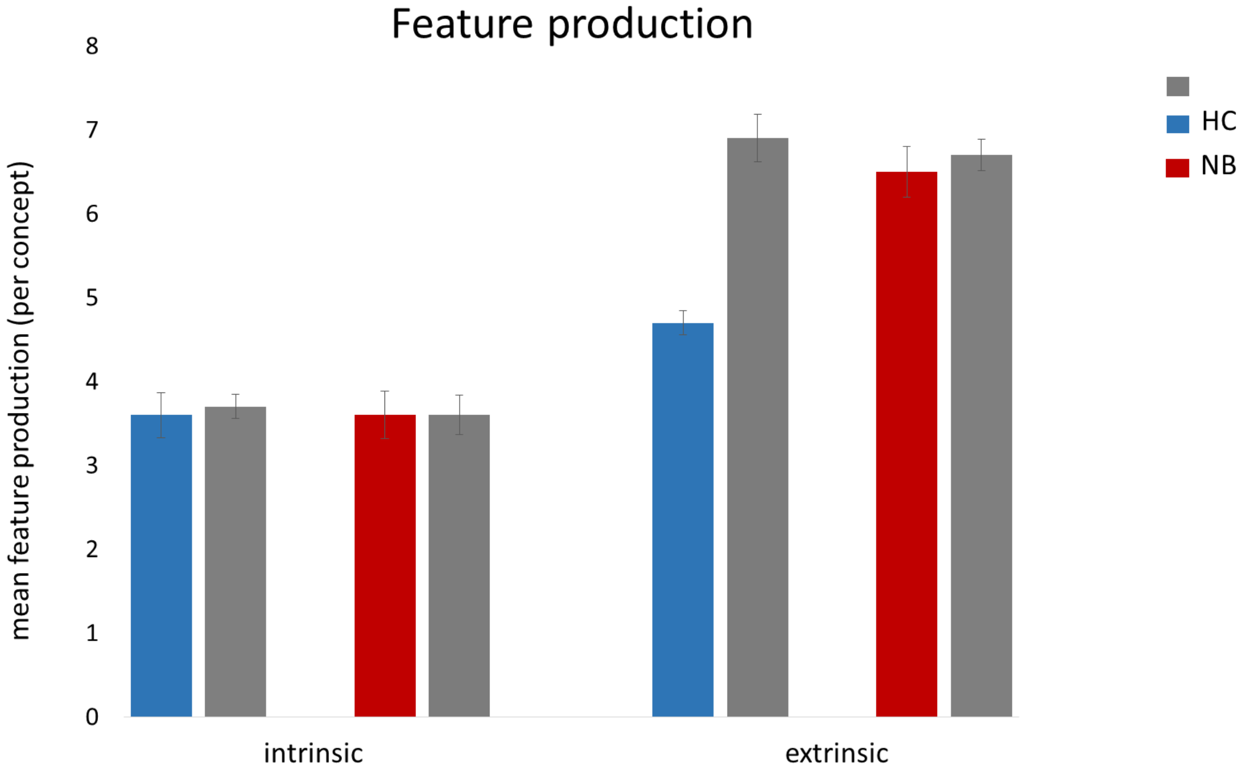

Figure 4.2: Mean intrinsic and extrinsic feature production for HC, NB, and their control participants. HC did not differ from controls in the number of intrinsic features produced, but produced significantly fewer extrinsic features (F(1,153)() =

44.54, p<0.0001). NB did not differ from controls in the number of features produced for intrinsic or extrinsic features.

4.4 Experiment 2: Typicality ratings

4.4.1 Materials

The stimuli were 429 basic-level concepts taken from 10 nonliving (e.g., *clothing-shirt*) and 7 living (e.g., *fruit-apple*) superordinate categories. The categories and their exemplars were based on the taxonomic features produced by participants in McRae et al.'s (2005) feature production database. That is, as part of producing features, participants in McRae et al. provided the category to which basic-level concepts belong, and those responses were used to determine the superordinate categories and their exemplars. The number of exemplars per superordinate category ranged from 9 to 57, with a mean of 25 (see Table 2). Some concepts were included as exemplars of multiple categories (e.g., *tractor* was included both for the category *vehicle* and for *machine*). We excluded ambiguous concepts (such as *bat*) and those that do not fit cleanly into any category (e.g., some food items).

Table 4.2 Number of exemplars used in each object category in the typicality rating experiment

Category	Number of Exemplars	Example Exemplars		
mammal	57	bat	elk	bear
bird	39	blackbird	flamingo	hawk
insect	13	hornet	ant	caterpillar
reptile	10	salamander	toad	iguana
carnivore	19	eagle	fox	hyena
fruit	29	apple	tomato	cranberry
vegetable	31	olive	zucchini	cauliflower
appliance	14	toaster	microwave	blender
utensil	22	mug	paintbrush	bowl
tool	34	axe	level	microscope
weapon	39	whip	sword	shield
container	14	ashtry	box	urn
furniture	15	bed	couch	stool
clothing	39	shirt	veil	trousers

machine	9	tank	catapult	helicopter
vehicle	27	yacht	bus	ship
musical	18	accordion	harp	trumpet

4.4.2 Procedure

For each trial, participants were presented with a superordinate category label near the top of the computer screen, and an exemplar name below. They were asked to use a 9-point scale to indicate how good of an example each basic-level concept is of the category listed above it, with ‘9’ corresponding to an extremely good member of a category, and ‘1’ to a poor category member. Participants rated all exemplars for a given category in random order, and then proceeded to the next category. The 17 categories also appeared in random order. All participants were instructed to pay close attention to subtle differences regarding the typicality dimension, and they were encouraged to use the entire scale. Ratings were self-paced and trials were separated by an ISI of 750 ms. Practice items to familiarize participants with the task used sports as the category. The experiment took approximately 40 minutes.

4.4.3 Results

We first assessed HC and NB’s performance in general by using analyses that did not distinguish between living and nonliving categories. We used either z-scores or a modified t-test that was developed specifically for experimental single-case studies (Crawford & Howell, 1998). HC’s overall mean typicality rating (6.28) did not differ significantly from that of controls (6.70), $z = -0.53$, $p > .59$. The standard deviation of HC’s (2.65) and controls’ (2.46) ratings also did not differ significantly, $z = 0.58$, $p > .55$.

NB did not differ in terms of her average typicality rating (6.54) from that of controls (6.85), $z = -0.33$, $p > .7$. The standard deviation of NB's (2.20) and controls' (2.25) typicality ratings also did not differ significantly, $z = -0.21$, $p > .8$. Thus, overall across all categories, neither HC nor NB provided significantly lower or higher typicality ratings than did controls.

Our primary interest was to determine whether HC rates typicality differently for nonliving as compared to living thing categories. Specifically, we hypothesized that because HC has less well-developed extrinsic feature knowledge, her ratings would be abnormal when compared to controls, although we did not expect her to generally find nonliving items more or less typical than controls. Therefore, we evaluated both the pattern of typicality ratings across living and nonliving items (using correlations), as well as HC's average typicality ratings for both categories. We used the correlational approach because we wanted to examine the pattern of typicality ratings, rather than their absolute magnitudes. In these ratings on a 1 to 9 scale, the absolute magnitude of the ratings is less important than are the relative ratings of typicality across the exemplars for each superordinate category. Our correlational approach was designed specifically to be sensitive to how participants rate various exemplars relative to one another. Given that NB did not differ from controls in terms of extrinsic feature knowledge in Experiment 1, we predicted that her pattern of typicality ratings would be similar to control participants for both living and nonliving things. To compare HC and NB's typicality ratings to those of controls, we calculated item-based correlations between each individual's typicality ratings and their controls' typicality ratings. First, we correlated HC's and NB's ratings with each of the controls and then calculated the average correlation. For each control, we correlated their ratings with those of every other control participant, and then calculated an average correlation. As depicted in Figure 4.3, for all 17 superordinate categories combined, the mean correlation between HC's typicality ratings and those of the controls was numerically but not significantly lower than the mean correlation of each control with the other controls (control mean: $r = .45$; HC: $r = .37$; $t(9) = -1.19$, $p > .26$). NB's overall correlation was similar to controls (control mean: $r=0.48$, NB mean: $r=0.50$; $t(10)=0.621$, p > .55).

The theoretically central analyses involved dividing the categories into living and nonliving things. HC's average correlation for living thing categories did not differ from those of controls (control mean: r = .33; HC: r = .30; t(9) = -0.26, p > .80). Importantly however, for nonliving thing categories, HC’s average correlation was the lowest, and was significantly lower than that of controls (control mean: r = .50; HC: r = .40; t(9) = -2.33, p < .05). Furthermore, HC did not differ with respect to the mean typicality rating given to nonliving (control mean = 6.33; HC = 5.94; z = -0.57, p > .56) nor living things (control mean = 7.13; HC = 6.68; z = -0.44, p > .66). In fact, HC tended to provide slightly lower typicality ratings for both types of concepts (0.39 lower for nonliving things, 0.45 lower for living things). This demonstrates that abnormalities in typicality ratings for nonliving concepts cannot be explained by HC simply rating all concepts generally as high or low in typicality. Rather, it is the pattern of typicality across the nonliving things that is abnormal in HC.

Unlike HC, NB’s pattern of typicality ratings did not differ from controls for nonliving things (control mean: r = 0.54, NB: r = 0.53, t(10) = -0.34, p > .74). Also, like HC, NB’s mean correlation did not differ for living thing categories (control mean: r = 0.37, NB: r = 0.44, t(10) = 0.92, p > .38). Finally, NB’s average typicality rating for the nonliving categories did not differ from controls (control mean = 6.55, NB mean = 6.28, t(10) = -0.10, p > .92), and the same was true for living things (control mean = 7.23, NB mean = 6.9, t(10) = -0.14, p > .90).

An important question that arises is whether or not there are factors that differ in general between living and nonliving concepts that could account for HC’s abnormal typicality ratings aside from an explanation based on extrinsic features. For example, there could be differences in terms of the frequency that nonliving and living concepts appear in text, or the concreteness of the two domains. To explore these possibilities, we evaluated frequency, as measured by *ln*(BNC), the natural logarithm of the number of times each word occurs in the British National Corpus. Across all 541 concepts in McRae et al.’s (2005) norms, frequency was significantly higher for nonliving things (mean living = 5.8, SD = 1.6; mean nonliving = 6.7, SD = 1.8; t(539) = 5.86, p < .001). We also analyzed concreteness ratings obtained from Brysbaert, Warriner, and Kuperman (2014) for the

538 concepts from McRae et al.'s norms for which concreteness ratings were available. Concreteness, on a 1 to 5 scale with 5 being the most concrete, did not differ between living and nonliving things (mean living = 4.8, SD = 0.2; mean nonliving = 4.8, SD = 0.2; $t(536) = 1.37$, $p > .16$). As would be expected, concreteness ratings for the living and nonliving things were uniformly extremely high.

Furthermore, we were able to test more sensitively the possibility that HC's difference in typicality ratings between nonliving and living things could be due to the possibility that HC is simply more familiar overall with living things. To test this, we used HC's familiarity ratings for all 541 concepts from McRae et al. (2005). These ratings were collected as part of another project (Bowles et al., 2016). HC's familiarity ratings did not differ from controls' ratings for nonliving (controls: $M = 5.91$, $SD = 1.00$; HC: $M = 4.83$, $z = - 1.08$, $p > .2$) or for living things (controls: $M = 4.55$, $SD = 1.37$; HC: $M = 3.30$, $z = 0.92$, $p > .3$). Although HC's familiarity ratings are numerically lower than control participants', if anything, the difference is larger for living than for nonliving things. *More importantly, HC rated the nonliving things as being more familiar to her than the living things*. It is also notable that in the literature on acquired category-specific deficits with adult onset, approximately 80% of reported patients have shown living thing deficits (Capitani, et al., 2003). Furthermore, prior to the use of improved designs in which familiarity was balanced across living and nonliving things, it had been argued that the preponderance of patients with living things deficits was due, at least in part, to living things being less familiar than nonliving things, rather than the other way around (Funnell & Sheridan, 1992).

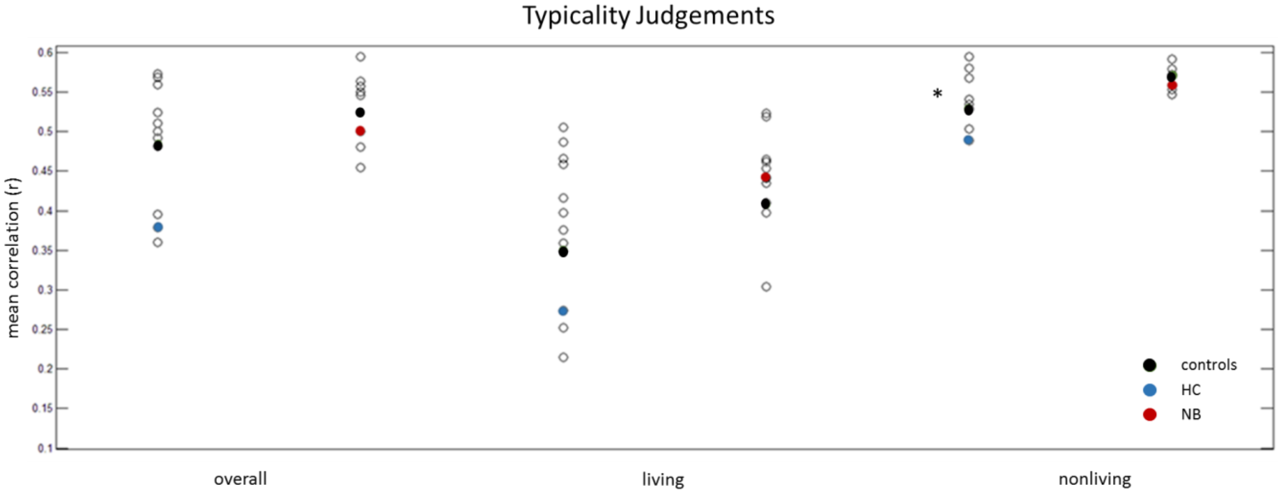

Figure 4.3 Typicality judgements for living and nonliving object concepts for HC, NB, and their control participants. Separate control groups were chosen to best match each individual, for the highest level of control. HC's typicality ratings do not significantly differ from controls overall (control mean: *r* = 0.45; HC: *r* = 0.37; *t*(9) = −1.19, p>0.26), or in the living domain (control mean: *r* = 0.33; HC: *r*= 0.30; *t*(9) = −0.26, p>0.80), but were significantly less correlated with control's ratings in the nonliving domain (control mean: *r* = 0.50; HC: *r* = 0.40; *t*(9) = −2.33, p<0.05). NB's typicality ratings did not differ from controls (overall: control mean: *r*=0.48, NB mean: *r* =0.50; *t*(10) =0.621, p>0.55; living: control mean: *r* = 0.37, NB: *r* = 0.44, *t*(10) = 0.92, p>0.38; nonliving: control mean: *r* = 0.54, NB: *r* = 0.53, *t*(10) = −0.34, p>0.74).

4.5 Discussion

Consistent with our main hypothesis, we found that HC, an individual with developmental amnesia, demonstrates some abnormalities in conceptual structure. These abnormalities appear to be specific to extrinsic features of object concepts. In Experiment 1, when asked to generate individual features for concrete concepts, HC produced significantly fewer extrinsic features than control participants, but was able to generate as many intrinsic features as controls. In Experiment 2, we used a typicality-rating task to test whether HC's impoverished extrinsic feature knowledge would lead to abnormal ratings for nonliving object concepts. When asked to rate the typicality of concepts in relation to a superordinate category, HC's ratings for nonliving concepts were significantly less correlated with control participants' ratings than the control participants' ratings were correlated with one another. Importantly, HC's typicality ratings for living things did not show any such abnormalities. Critically, these results cannot be accounted for by reduced familiarity specific to nonliving objects (weapons, tools, etc.). When we examined HC's lifetime familiarity ratings for all concepts used in our experiment, there were no significant differences between her ratings and those of control participants for both living and nonliving concepts. Additionally, NB, an individual with medial temporal lobe damage sparing the hippocampus, performed similarly to controls on both the feature generation and typicality tasks, suggesting the hippocampus specifically is important for extrinsic feature knowledge. Although caution is required when comparing patients with different etiologies, these findings provide some specificity as to the lesion location that may cause abnormalities in extrinsic feature knowledge. However, inasmuch as NB's lesion was acquired later in life than HC's, further research is required in other patients to shed light on this issue.

As is typically the case in individuals with developmental amnesia (Vargha-Khadem et al., 1997), many aspects of semantic knowledge appear to be spared in HC. She performs normally on standard semantic knowledge tests, such as object naming, category fluency with animals, fact-based questionnaires about the world, and vocabulary tests (Table 1, see for further detail Rosenbaum et al., 2011). Such findings have been taken to support a strong claim that semantic memory acquisition precedes and is independent of hippocampally-based episodic memory (Tulving and Markowitsch, 1998; Tulving, 2002; Vargha-Khadem et al., 1997; de-Haan et al., 2006). However, the clinical neuropsychological tests of semantic memory that were used to establish preserved semantic memory were not designed to probe specific aspects of conceptual structure. The experimental tasks used in the current study, by contrast, allowed us to test specific hypotheses concerning the structure of concrete concepts. Experiment 1 allowed us to compare the generation of extrinsic and intrinsic features, and revealed an abnormality in

HC’s performance that is specific to the former. In Experiment 2, typicality ratings allowed us to show, building on a longstanding research strategy in the semantic memory and categorization literatures (Rosch & Mervis, 1975), that this abnormality in knowledge of extrinsic features has implications for the internal structure of some but not all categories and their exemplars.

To the extent that the pathology in developmental amnesia centers around the hippocampus, and in HC includes the extended hippocampal system, extrinsic features of concepts are of particular relevance. Prominent theories of hippocampal functioning emphasize its role in relational binding, or the binding of disparate elements (such as objects and contextual information) across space and time. Traditionally, this binding function has been discussed in terms of its role in episodic memory (Davachi, 2006; Mayes et al., 2007; Konkel & Cohen, 2008; Cohen & Eichenbaum, 1993). However, our results suggest that hippocampal binding may also be important for certain aspects of semantic memory.

The hippocampus is critical for binding objects with associated extrinsic or contextual information, such as the background scene or the specific list in which the item was initially presented (Mitchell & Johnson, 2009). When contextual information is external to the object itself, patients with hippocampal lesions are impaired in recalling the source (Ryan et al., 2000; Hannula et al., 2007). Critically, however, when individuals with hippocampal lesions are instructed to “unitize” an item and its context (e.g., imagine the object as having the same colour as the background, making the context part of the item), they are able to retrieve the unitized associations (Diana et al., 2009; cf. Ryan et al., 2013; D’Angelo et al., 2015). Functional neuroimaging results demonstrate that when items are unitized during encoding, the perirhinal cortex, and not the hippocampus, is activated (Diana et al., 2009; Haskins et al., 2008). HC’s intact perirhinal cortex (Olsen et al., 2013) may therefore explain why she is unimpaired in producing intrinsic features, and normal in typicality measures for living things.

The reported abnormalities following early hippocampal damage speak to the notion that at least some aspects of semantic memory are a result of a transformation process by which information occurring within unique episodic experiences becomes more abstract and schematized over time as part of consolidation (Moscovitch et al., 2016; Winocur et al., 2010). This transformation process is thought to be supported by interactions between the hippocampus and ventromedial prefrontal cortex. The described abnormalities for extrinsic feature knowledge in H.C. can be interpreted as the outcome of impairments in this process, due to poorly bound episodic experiences related to pertinent concepts in initial hippocampal representations. As such, the present results on extrinsic feature knowledge are not easily accommodated by the SPI model of semantic and episodic memory (Tulving & Markowitsch, 1998), in which the development of semantic memory representations always precedes, and is necessary for, encoding in episodic memory. Thus, while the SPI model may account for the intact acquisition of some aspects of semantic knowledge in developmental amnesia, it appears that it does not hold for semantic learning universally.

Other evidence linking semantic learning to hippocampal processing comes from functional neuroimaging research that has revealed a role for the hippocampus in binding information across multiple episodes. In a study by Kumaran et al. (2009), participants had to learn to predict outcomes (rain or shine) for patterns composed of two fractals based on either object-object pairings, or object-location pairings. The task was structured so that participants could extract diagnostic information across trials based on some commonality, against the background of varying irrelevant information. Accumulation of this conceptual knowledge was related to activity in a functionally coupled circuit between the hippocampus and ventromedial prefrontal cortex. The binding of outcome information (rain or shine) to a pattern (fractal combination) may be thought of as a process akin to extrinsic feature binding of an object (scissors) to functional information about the object, such as its purpose (to cut). Developing a rich and stable concept requires binding an object with important extrinsic features which repeatedly co-occur across multiple episodes, or are particularly salient (i.e., the act of cutting).

The notion that learning semantic knowledge that requires feature binding relies on the hippocampus also receives support from findings in individuals with adult-onset amnesia.

While a detailed consideration of this literature is beyond the scope of the current paper, we note that many studies have reported abnormalities in learning new semantic information. Holdstock et al. (2002) evaluated an individual, YR, on her ability to rapidly learn new semantic information. YR suffered an ischaemic accident leaving her with selective bilateral hippocampal lesions and anterograde amnesia, her neuropsychological profile shows she has normal intelligence, and intact recognition of visual and verbal items, with a fairly specific impairment in recall. In the current study, YR was tested on her ability to learn a set of new word-definition pairs, which were shown on index cards and read out loud in ten different sessions. YR's recall was impaired relative to controls on testing immediately after each session, and after a 30 minute delay period, indicating that she was unable to rapidly learn new vocabulary words. The authors interpreted YR's deficit in acquiring new semantic knowledge as an inability to bind orthographic information to its meaning, suggesting that the same hippocampal mechanisms that bind elements within an episode are also needed to acquire new semantic knowledge (Holdstock et al., 2000). In other studies on individuals with adult-onset amnesia, some acquisition of semantic knowledge has been reported, but typically at a slower rate than in control participants (Bayley & Squire, 2002; Verfaellie et al., 2000; O'Kane et al., 2004; Tulving, Hayman, & MacDonald, 1991; but see Sharon et al., 2011). It has been suggested that this slow acquisition relies on cortical learning mechanisms that are distinct from relational-binding mechanisms supported by the hippocampus, perhaps including but likely not restricted to computations in perirhinal cortex and the anterior temporal lobes (Norman & O'Reilly, 2003). Adult amnesics may also be capable of rapid incidental concept learning when information is actively discovered by inference based on previous knowledge ('fast-mapping'; Sharon et al., 2011; Merhav et al., 2014; but see Smith et al., 2014). Imaging and lesions studies point to the anterior temporal lobes as possible critical nodes for such learning (Merhav et al., 2015; Atir-Sharon et al., 2015). Fast-mapping is thought to be responsible for the prodigious rate at which children acquire vocabulary, and if it is cortically mediated independently of the hippocampus, that could account for preserved basic conceptual learning in developmental amnesia (as proposed by the SPI model discussed previously). Fast-mapping, however, only enables acquisition of partial meaning of a concept, and deeper more extensive conceptual knowledge may only be acquired more gradually through multiple experiences.

Interestingly, some recent evidence also shows that individuals with adult-onset amnesia due to hippocampal damage perform worse on measures of vocabulary depth and semantic richness even for premorbidly acquired concepts. While these individuals were not impaired in naming tasks and in tasks requiring matching a word to its definition, aspects of semantic knowledge which are thought to accumulate continuously over time and experience, such as vocabulary depth and semantic richness, were impoverished (Klooster et al., 2015; cf. Rosenbaum et al., 2009). The measure of semantic richness found to be affected in these patients required listing as many features as possible for a subset of concepts taken from McRae et al. (2005). In this task, individuals with adult onset amnesia produced fewer features overall, although the authors did not classify features into extrinsic and intrinsic domains as we did in the current study. One possible interpretation suggested by the authors is that the hippocampus is needed for incorporating or binding new features, even to a familiar concept, across repeated episodes. An important avenue for future research is to test whether incorporating new features is impoverished primarily for extrinsic features in adult-onset amnesia.

Dissociations between knowledge for living and nonliving things in adult neuropsychological patients has been widely reported. In general, category-specific semantic deficits have been reported more frequently (approximately 80% of cases) for living than for nonliving things. However, there are some cases of deficits specific to the nonliving domain. A recent review by Capitani et al. (2003) describes 18 case studies in which individuals were reported with a primary nonliving thing (artefact) impairment (p. 239). Of these 18 reported cases, 6 reported damage in the temporal lobes. One interesting possibility is that some of these cases may have involved damage to the hippocampus. Of these 6 cases, 5 involved damage primarily to the left temporal lobe (Warrington & McCarthy, 1987; Lambon Ralph et al., 1998; Cappa et al., 1998; Silveri et al., 1997; Moss & Tyler, 2000), however due to the fact that most lesions were documented by CT scans, or the damage was fairly widespread within the left temporal lobe, it is difficult to determine with any confidence whether the nonliving deficits reported in prior cases result from hippocampal damage. It seems to be an open question as to whether adult-onset amnesics with damage to the hippocampus show any abnormalities in extrinsic feature knowledge or processing nonliving objects, or as to whether the hippocampus is crucial for acquiring, but not retrieving/using extrinsic feature knowledge.

There is some evidence from neuropsychological research that acquiring extrinsic features, specifically functional knowledge, is impaired in adult onset amnesia. In a novel tool use paradigm, amnesic individual DA showed impaired learning of functional information related to novel tools (Roy & Park, 2010). DA was able to learn the series of movements required to use the tools as quickly as control participants across a training session, and he retained these motor skills overtime, as evidenced during later testing. However, even after extensive exposures across multiple sessions, DA could not recall tool attributes, including their function. In addition to this verbal recall impairment for functional information, DA showed abnormalities on two measures of tool use that required expression of this newly learned functional semantic knowledge, in a grasp-to-command and a use-to-command task. Notably, DA was also impaired in recalling the color of the tools, which is a perceptual intrinsic feature. This result contrasts with the preserved intrinsic feature knowledge we report here for HC. We also note however, that the test format used in this study involved a procedure that did not present color as an intrinsic feature of the tool, which may have limited demonstration of preserved perceptual knowledge. Specifically, a colored square separate from the tool itself was used to probe memory for tool color, i.e., acquired perceptual feature knowledge. It is possible that DA would perform better in a forced choice test in which color was presented as an intrinsic feature, as it would have been during learning.

Additional support for the notion that functional information relies on the hippocampus comes from recent fMRI research in healthy participants. In Chen et al. (2015), participants were presented with names of commonly used tools, and were required to pantomime the tool’s corresponding action. Motor-actions for different types of tools could be successfully classified based on representations in dorsal and ventral stream areas in the parietal and temporal lobes, whereas functional information about tools could be classified almost exclusively through activity patterns in the medial temporal lobes, including in the hippocampus (Chen et al., 2015). While this particular study did not examine the representation of intrinsic features, it adds to a body of fMRI studies that provide evidence for unique representations of functional knowledge as compared to other aspects of object knowledge (Mahon & Caramazza, 2011).

On another level, the finding that medial temporal lobe structures represent tool function information during action pantomiming adds to a broader emerging literature that implicates the medial temporal lobes in tasks that require expression of semantic knowledge. For example, recent research has provided evidence that the hippocampus makes contributions to tasks such as object naming or conceptual fluency, that is, the speeded generation of exemplars from different semantic categories (Greenberg et al., 2009; Ryan et al., 2008; Sheldon et al., 2012; Westmacott & Moscovitch, 2003; Whatmough & Chertkow, 2007). Building again on the view that this structure plays a critical role in binding items to episodic contexts (Cohen and Eichenbaum, 1993), such evidence has also led to the suggestion that hippocampal representations support cognition well beyond their widely accepted role in episodic memory (see Moscovitch et al., 2015, for detailed discussion). The conclusions of this broader literature converge with the findings on developmental amnesia reported here, in questioning a sharp distinction between episodic and semantic memory in functional brain organization.

4.5.1 Conclusion

The current case study in developmental amnesia with HC and NB reveals an important role for the hippocampus in learning extrinsic semantic features of object concepts. While some aspects of semantics are intact in HC, consistent with the SPI model (Tulving and Markowitsch, 1998; Tulving, 2002) and prior experimental reports (Vargha-Khadem et al., 1997; Vargha-Khadem et al., 2003; Baddeley et al., 2001; Martins et al., 2006), our novel findings emphasize that the development of important aspects of conceptual knowledge might still depend on the hippocampus. Our results demonstrate that performance on semantic memory tasks that tap into the feature structure of concrete concepts can be affected by early hippocampal pathology for concepts that are characterized by high levels of extrinsic associative information. This role in extrinsic feature binding is consistent with the classic notion that the hippocampus binds objects to contexts. Although this notion is well established in the episodic memory literature, the current findings suggest that it also has implications for the representation of concepts in semantic memory.

4.6 References

Adlam, A. L. R., Vargha-Khadem, F., Mishkin, M., & De Haan, M. (2005). Deferred imitation of action sequences in developmental amnesia. *Journal of Cognitive Neuroscience*, *17*(2), 240-248.

Adlam, A. L. R., Malloy, M., Mishkin, M., Vargha-Khadem, F. (2009). Dissociation between recognition and recall in developmental amnesia. *Neuropsychologia*, *47*(11), 2207-2210.

Atir-Sharon, T., Gilboa, A., Hazan, H., Koilis, E., & Manevitz, L. M. (2015). Decoding the formation of new semantics: MVPA investigation of rapid neocortical plasticity during associative encoding through Fast Mapping. *Neural plasticity, 2015*.

Baddeley, A., Vargha-Khadem, F., & Mishkin, M. (2001). Preserved Recognition in a Case of Developmental Amnesia: Implications for the Acquisition of Semantic Memory? *Journal of Cognitive Neuroscience*, *13*(3), 357-369.

Barr, R. A. & Caplan, L. J. (1987). Category representations and their implications for category structure. *Memory & Cognition*, *15*, 397-418.

Barton, M. E. & Komatsu, L. K. (1989). Defining features of natural kinds and artifacts. *Journal of Psycholinguistic Research*, *18*, 433-447.

Bayley, P. J., & Squire, L. R. (2002). Medial temporal lobe amnesia: Gradual acquisition of factual information by nondeclarative memory. *The Journal of Neuroscience*, *22*(13), 5741-5748.

Bindschaedler, C., Peter-Favre, C., Maeder, P., Hirsbrunner, T. & Clarke, S. (2011). Growing up with bilateral hippocampal atrophy: From childhood to teenage. *Cortex*, *47*, 931-944.

Bowles, B., Duke, D., Rosenbaum, R. S., McRae, K. & Köhler, S. Impaired assessment of cumulative lifetime familiarity for object concepts after left anterior temporal-lobe resection that includes perirhinal cortex but spares the hippocampus. *Neuropsychologia*, *90*, 170-179.

Brizzolara, D., Casalini, C., Montanaro, D. & Posteraro, F. (2003). A case of amnesia at an early age. *Cortex*, *39*, 605-625.

Brysbaert, M., Warriner, A. B, & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, *46*, 904-911.

Capitani, E., Laiacona, M., Mahon, B. Z., & Caramazza, A. (2003). What are the facts of semantic category-specific deficits: A critical review of the clinical evidence. *Cognitive Neuropsychology*, *20*, 213-261.

Cappa, S. F., Frugoni, M., Pasquali, P., Perani, D., & Zorat, F. (1998). Category-specific naming impairment for artefacts: A new case. *Neurocase*, *4*(4-5), 391-397.

Casler, K., & Kelemen, D. (2007). Reasoning about artifacts at 24 months: The developing teleo-functional stance. *Cognition*, *103*, 120-130.

Chen, Q., Garcea, F. E. & Mahon, B. Z. (2015). The representation of object-directed action and function knowledge in the human brain. *Cerebral Cortex, 1-10*.

Cohen, N. J. & Eichenbaum, H. (1993). Memory, amnesia, and the hippocampal system. MIT Press, Cambridge, MA.

Cree, G. S. & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of *chipmunk*, *cherry*, *chisel*, *cheese,* and *cello (*and many other such concrete nouns). *Journal of Experimental Psychology: General*, *132*(2), 163-201.

Crawford, J. R. & Howell, D. C. (1998). Comparing and individualès test score against norms derived from small samples. *The Clinical Neuropsychologist*, *12*(4), 482-486.

D'Angelo, M. C., Kacollja, A., Rabin, J. S., Rosenbaum, R. S., & Ryan, J. D. (2015). Unitization supports lasting performance and generalization on a relational memory task: Evidence from a previously undocumented developmental amnesic case. *Neuropsychologia*, *77*, 185-200.

Davachi, L. (2006). Item, context and relational episodic encoding in humans. *Current opinion in neurobiology*, *16*(6), 693-700.

De Haan, M., Wyatt, J. S., Roth, S., Vargha-Khadem, F., Gadian, D., & Mishkin, M. (2006). Brain and cognitive-behavioural development after asphyxia at term birth. *Developmental science*, *9*(4), 350-358.

Diana, R. A., Yonelinas, A. P. & Ranganath, C. (2009). Medial temporal lobe activity during source retrieval reflects information type, not memory strength. *Journal of Cognitive Neuroscience, 22*(8), 1808-1818.

Diesendruck, G., Hammer, R., & Catz, O. (2003). Mapping the similarity space of children and adults' artifact categories. *Cognitive Development*, *18*, 217-231.

Farah, M. J. & McClelland, J. L. (1991). A computational model of semantic memory impairment: Modality specificity and emergent category specificity. *Journal of Experimental Psychology: General, 120(4),* 339-357.

Funnell, E. & Sheridan, J. (1992). Categories of knowledge? Unfamiliar aspects of living and nonliving things. *Cognitive Neuropsychology*, *9*, 135-153.

Gabrieli, J. D., Cohen, N. J., & Corkin, S. (1988). The impaired learning of semantic knowledge following bilateral medial temporal-lobe resection. *Brain and cognition*, *7*(2), 157-177.

Gadian, D. G., Aicardi, J., Watkins, K. E., Porter, D. A., Mishkin, M. & Vargha-Khadem, F. (2000). Developmental amnesia associated with early hypoxic-ischaemic injury. *Brain, 2000, 123*, 499-507.

Gardiner, J. M., Brandt, K. R., Baddeley, A.D., Vargha-Khadem, F. & Mishkin, M. (2008). Charting the acquisition of semantic knowledge in a case of developmental amnesia. *Neuropsychologia,46*, 2865-2868.

Greenberg, D. L., Keane, M. M., Ryan, L., & Verfaellie, M. (2009). Impaired category fluency in medial temporal lobe amnesia: The role of episodic memory. *Journal of Neuroscience, 29* (35),10900–10908.

Guillery-Girard, B., Martins, S., Parisot-Carbuccia, D. & Eustache, F. (2004). Semantic acquisition in childhood amnesic syndrome: a prospective study. *NeuroReport, 15*(2), 377-381.

Hampton, J. A., Storms, G., Simmons, C. L., & Heussen, D. (2009). Feature integration in natural language concepts. *Memory & Cognition*, *37*, 1150-1163.

Hannula, D. E., Tranel, D., & Cohen, N. J. (2006). The long and short of it: Relational memory impairments in amnesia, even at short lags. *Journal of Neuroscience, 26*, 8352-8359.

Haskins, A. L., Yonelinas, A. P., Quamme, J. R., & Ranganath, C. (2008). Perirhinal cortex supports encoding and familiarity-based recognition of novel associations. *Neuron, 59*(4), 554-560.

Holdstock, J. R., Mayes, A. R., Isaac, C. L. Gong, Q., & Roberts, N. (2002). Differential involvement of the hippocampus and temporal lobe cortices in rapid and slow learning of new semantic information. *Neuropsychologia*, *40*, 748-768.

Kalénine, S., & Bonthoux, F. (2008). Object manipulability affects children’s and adults’ conceptual processing. *Psychonomic Bulletin & Review*, *15*, 667–672.

Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.

Kemler Nelson, D. G., Chan Egan, L., & Holt, M. B. (2004). When children ask, “what is it?” What do they want to know about artifacts? *Psychological Science*, *15*, 384-389.

Kinsbourne, M. & Wood, F. Short term memory processes and the amnesic syndrome. D. Deutsch, J.A. Deutsch (Eds.), Short-term memory, Academic Press, New York (1975), pp. 258–291

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*.

Klooster, N. B. & Duff, M. C. (2015). Remote semantic memory is impoverished in hippocampal amnesia. *Neuropsychologia, 79,* 45-52.

Kumaran, D., Summerfield, J. J., Hassabis, D. & Maguire, E. A. (2009). Tracking the emergence of conceptual knowledge during human decision making. *Neuron, 63,* 889-901.

Kwan, D., Carson, N., Addis, D. R., & Rosenbaum, R. S. (2010). Deficits in past remembering extend to future imagining in a case of developmental amnesia. *Neuropsychologia, 48*(11), 3179-3186.

Lavenex, P. & Amaral, D. G. (2000). Hippocampal-neocortical interaction: A hierarchy of associativity. *Hippocampus, 10*, 420-430.

Lambon Ralph, M. A., Howard, D., Nightingale, G., & Ellis, A. W. (1998). Are living and non-living category-specific deficits causally linked to impaired perceptual or associative knowledge? Evidence from a category-specific double dissociation. *Neurocase*, *4*(4-5), 311-338.

Mahon, B. Z. & Caramazza, A. (2011). What drives the organization of object knowledge in the brain? *Trends in Cognitive Science, 15*(3), 97-103.

Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology, 58,* 25-45.

Martins, S., Guillery-Girard, B., Jambaqué, I., Dulac, O., & Eustache, F. (2006). How children suffering severe amnesic syndrome acquire new concepts? *Neuropsychologia, 44*, 2792–2805.

Mayes, A., Montaldi, D., & Migo, E. (2007). Associative memory and the medial temporal lobes. *Trends in cognitive sciences*, *11*(3), 126-135.

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.

McRae, K., Cree, G. S., Seidenberg, M. S., & McNorgan, C. (2005). Semantic feature production norms for a large set of living and nonliving things. *Behavior research methods*, *37*(4), 547-559.

Medin, D. L. & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85*(3), 207-238.

Merhav, M., Karni, A., & Gilboa, A. (2014). Neocortical catastrophic interference in healthy and amnesic adults: a paradoxical matter of time. *Hippocampus*, *24*(12), 1653-1662.

Merhav, M., Karni, A., & Gilboa, A. (2015). Not all declarative memories are created equal: fast mapping as a direct route to cortical declarative representations. *NeuroImage*, *117*, 80-92.

Mitchell, K. J., & Johnson, M. K. (2009). Source monitoring 15 years later: What have we learned from fMRI about the neural mechanisms of source memory. *Psychological Bulletin*, *135*(4), 638-677.

Mishkin, M. (1997). Differential effects of early hippocampal pathology on episodic and semantic memory. *Science*, *277*(5324), 376-380.

Moss, H. E., & Tyler, L. K. (1997, October). A category-specific semantic deficit for nonliving things in a case of progressive aphasia. In *Brain and Language* (Vol. 60, No. 1, pp. 55-58).

Moss, H. E., & Tyler, L. K. (2000). A progressive category-specific semantic deficit for non-living things. *Neuropsychologia*, *38*(1), 60-82.

Nadel, L. & Peterson, M. A. (2013). The hippocampus: part of an interactive posterior representational system spanning perceptual and memory systems. *Journal of Experimental Psychology: General, 142*, 1242-1254.

Norman, K. A. & O'Reilly, R. C. (2003). Modeling hippocampal and neocortical contributions to recognition memory: A complementary-learning-systems approach. *Psychological Review, 110*(4), 611-646.

O'Kane, G., Kensinger, E. A., & Corkin, S. (2004). Evidence for semantic learning in profound amnesia: an investigation with patient HM. *Hippocampus,14*(4), 417-425.

Olsen, R. K., Palombo, D. J., Rabin, J. S., Levine, B., Ryan, J. D., & Rosenbaum, R. S. (2013). Volumetric analysis of medial temporal lobe subregions in developmental amnesia using high-resolution magnetic resonance imaging. *Hippocampus*, *23*(10), 855-860.

Olsen, R. K., Moses, S. N., Riggs, L. & Ryan, J. D. (2012). The hippocampus supports multiple cognitive processes through relational binding and comparison. *Frontiers in Human Neuroscience, 6*(145), 1-13.

Parkin, A. J. (1982). Residual learning capability in organic amnesia. *Cortex, 18*(3), 417-440.

Picard, L., Mayor-Dubois, C., Maeder, P., Kalenzaga, S., Abram, M., Duval, C., et al. (2013). Functional independence within the self-memory system: New insights from two cases of developmental amnesia. *Cortex, 49*, 1463-1481.

Rabin, J. S., Braverman, A., Gilboa, A., Stuss, D. T., & Rosenbaum, R. S. (2012). Theory of mind development can withstand compromised episodic memory development. *Neuropsychologia*, *50*, 3781–3785.

Rosenbaum, R. S., Gilboa, A., Levine, B., Winocur, G., & Moscovitch, M. (2009). Amnesia as an impairment of detail generation and binding: evidence from personal, fictional, and semantic narratives in KC. *Neuropsychologia*, *47*(11), 2181-2187.

Rosenbaum, R. S., Gao, F., Honjo, K., Raybaud, C., Olsen, R. K., Palombo, D. J., ... & Black, S. E. (2015). Congenital absence of the mammillary bodies: A novel finding in a well-studied case of developmental amnesia. *Neuropsychologia*, *65*, 82-87.

Rosenbaum, R. S., Carson, N., Abraham, N., Bowles, B., Kwan, D., Köhler, et al. (2011). Impaired event memory and recollection in a case of developmental amnesia. *Neurocase, 5*, 394-409.

Rosch, E. & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology, 7*(4), 573-605.

Roy, S., & Park, N. W. (2010). Dissociating the memory systems mediating complex tool knowledge and skills. *Neuropsychologia*, *48*(10), 3026-3036.

Ryan, J. D., Althoff, R. R., Whitlow, S. & Cohen, N. J. (2000). Amnesia is a deficit in relational memory. *Psychological Science*, *11*, 454-461.

Ryan L., Cox, C., Hayes, S. M. & Nadel, L. (2008). Hippocampal activation during episodic and semantic memory retrieval: comparing category production and category cued recall. *Neuropsychologia, 46*(8), 2109-2121.

Ryan, J. D., Moses, S. N., Barense, M., & Rosenbaum, R. S. (2013). Intact learning of new relations in amnesia as achieved through unitization. *The Journal of Neuroscience*, *33*(23), 9601-9613.

Schacter, D. L., & Tulving, E. (1994). *Memory systems 1994*. Mit Press.

Sharon, T., Moscovitch, M., & Gilboa, A. (2011). Rapid neocortical acquisition of long-term arbitrary associations independent of the hippocampus. *Proceedings of the National Academy of Sciences*, *108*(3), 1146-1151.

Sheldon, S., & Moscovitch, M. (2012). The nature and time-course of medial temporal lobe contributions to semantic retrieval: An fMRI study on verbal fluency. *Hippocampus*, *22*(6), 1451-1466.

Shimamura, A. P., & Squire, L. R. (1987). A neuropsychological study of fact memory and source amnesia. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*(3), 464.

Silveri, M. C., Gainotti, G., Perani, D., Cappelletti, J. Y., Carbone, G., & Fazio, F. (1997). Naming deficit for non-living items: Neuropsychological and PET study. *Neuropsychologia*, *35*(3), 359-367.

Simmons, W. K. & Barsalou, L. W. (2003). The Similarity-In-Topography principle: Reconciling theories of conceptual deficits. *Cognitive Neuropsychology*, *20,* 451-486.

Sharon, T., Moscovitch, M., & Gilboa, A. (2011). Rapid neocortical acquisition of long-term arbitrary associations independent of the hippocampus. *Proceedings of the National Academy of Sciences*, *108*(3), 1146-1151.

Sloutsky, V. M., Kloos, H., & Fisher, A. V. (2007). When looks are everything: Appearance similarity versus kind information in early induction. *Psychological Science*, *18*, 179-185.

Sloutsky, V. M., & Spino, M. A. (2004). Naive theory and transfer of learning: When less is more and more is less. *Psychonomic Bulletin & Review*, *11*, 528-535.

Smith, C. N., Urgolites, Z. J., Hopkins, R. O., & Squire, L. R. (2014). Comparison of explicit and incidental learning strategies in memory-impaired patients. *Proceedings of the National Academy of Sciences*, *111*(1), 475-479.

Squire, L. R. (1987). Memory and brain.

Squire, L. R., & Zola, S. M. (1998). Episodic memory, semantic memory, and amnesia. *Hippocampus*, *8*(3), 205-211.

Truxaw, D., Krasnow, M. M., Woods, C., & German, T. P. (2006). Conditions under which function information attenuates name extension via shape. *Psychological Science*, *17*, 367-371.

Tulving, E. (2002). Episodic memory: From mind to brain. *Annual Review of Psychology*, *53*, 1-25.

Tulving, E., Hayman, C. A., & Macdonald, C. A. (1991). Long-lasting perceptual priming and semantic learning in amnesia: A case experiment. Journal of Experimental Psychology: Learning, Memory, and Cognition, 17, 595–617.

Tulving, E., & Markowitsch, H. J. (1998). Episodic and declarative memory: role of the hippocampus. *Hippocampus*, *8*(3), 198-204.

Tulving, E. (2002). Episodic memory: From mind to brain. *Annual Review of Psychology*, *53,* 1-25.

Tyler, L. K., Chiu, S., Zhuang, J., Randall, B., Devereux, B. J., Wright, P., Clarke, A., & Taylor,

K. I. (2013). Objects and categories: Feature statistics and object processing in the ventral

stream. *Journal of Cognitive Neuroscience*, *25*, 1723-1735.

Waidergoren, S., Segalowicz, J., & Gilboa, A. (2012). Semantic memory recognition is supported by intrinsic recollection-like processes:“The butcher on the bus” revisited. *Neuropsychologia*, *50*(14), 3573-3587.

Warrington, E. K., & Mccarthy, R. A. (2002). Categories of knowledge: further fractionations and an attempted integration. *Neurocase*, *8*(3), 192.

Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain, 107*(3), 829-853.

Westmacott, R., & Moscovitch, M. (2003). The contribution of autobiographical significance to semantic memory. *Memory and Cognition, 31*(5), 761-774.

Whatmough, C. & Chertkow, H. (2007). rCBF to the hippocampal complex covaries with superior memory retrieval. *Behavioral Brain Research, 181*(2), 262-269.

Winocur, G., Moscovitch, M., & Bontempi, B. (2010). Memory formation and long-term retention in humans and animals: Convergence towards a transformation account of hippocampal–neocortical interactions. *Neuropsychologia, 48*(8), 2339-2356.

Vargha-Khadem, F., Gadian, D. G., Watkins, K. E., Connelly, A., Van Paesschen, W. & Mishkin, M. (1997). Differential effects of early hippocampal pathology on episodic and semantic memory. *Science, 277,* 376-380.

Vargha-Khadem, F., Salmond, C. H., Watkins, K. E., Friston, K. J., Gadian, D. G. & Mishkin, M. (2003). Developmental amnesia: Effect of age at injury. *Proceedings of the National Academy of Sciences, 2003, 100*(17), 10055-10060.

Verfaellie, M., Koseff, P., & Alexander, M. P. (2000). Acquisition of novel semantic information in amnesia: effects of lesion location. *Neuropsychologia*, *38*(4), 484-492.

Vicari, S., Menghini, D., Di Paolo, M., Serra, L., Donfrancesco, A., Fidani, P., et al. (2007). Acquired amnesia in childhood: a single case study. *Neuropsychologia, 45,* 704-715.

Wu, L. L. & Barsalou, L. W. (2009) Perceptual simulation in conceptual combination: evidence from property generation. *Acta Psychologica, 132*(2), 173-189.

Zola-Morgan, S., Cohen, N. J., & Squire, L. R. (1983). Recall of remote episodic memory in amnesia. *Neuropsychologia*, *21*(5), 487-500.

Chapter 5

5 General Discussion

5.1 Object processing in the MTL: influence of object domain

A wealth of evidence suggests that MTL structures are involved in object perception, recognition memory for objects, and object concepts. One useful approach to understanding how MTL structures contribute to object processing is by focusing on different *types* of objects, in terms of category membership and domain. A focus on differentiation in representational content across MTL structures by category and domain can contribute to our understanding of the role these structures play more generally in behavioral and cognitive processes. I explored how object category, and more centrally object domain, drives the organization of responses evoked by visual object stimuli across MTL structures (Chapter 2), whether object domain and category also shape recognition-memory signals within these structures (Chapter 3), and finally how one MTL structure, the HpC, contributes to the acquisition of object concepts in a domain-specific manner (Chapter 4). In this General Discussion, I first review key findings from each chapter. I then discuss these findings and how they further our knowledge of MTL in the context of the broader theoretical literature, and make some suggestions for future research. I end with some broad conclusions and remaining open questions.

5.2 Organization of object representations related to stimulus processing in the MTL

An important aspect of understanding how MTL structures contribute to object perception and memory is understanding their sensitivity to content. While the previous literature suggested some differentiation in the types of stimuli represented by PhC and PrC, with the HpC generally being agnostic to stimulus type, prior to my study, this had not thoroughly been explored with a large number of object categories. Further, it was unclear whether these regions, like more posterior VVS regions, were organized by broader domains - in particular animacy and real-word size for inanimate objects. To address these questions, we scanned participants on a continuous recognition memory task with exemplars from a large number of object categories that could be divided by whether objects were animate or inanimate, and in the inanimate domain, whether they were large or small. We then mapped out the organization of representational space in each structure for stimulus processing (Chapter 2) and memory signals (Chapter 3).

In Chapter 2, we found that PhC represented the four large inanimate object categories tested (buildings, furniture, trees, and vehicles), as well as faces, and that representations were organized by both animacy and real-world size for inanimate objects. PrC was found to represent monkeys, with a trend for faces, and was organized more broadly by animacy but not real-world size. The HpC showed a trend for representing buildings, and was organized more broadly by real-world size for inanimate objects.

The focus of this fMRI study was to map out representational space for a large number of object categories, and to investigate whether animacy and real-world size are organizing dimensions within the MTL, as they are in much of visual cortex. Of course, this does not yet get at the question of *why* animacy and real-world size for inanimate objects are prominent dimensions at the level of anatomical and representational space. One possibility is that the MTL reflects the organization already present in upstream object processing regions that feed into these structures. Indeed, PrC, which we found to be organized by animacy but not real-world size, and which showed category organization for monkeys (with faces trending), has more dense innervations from the lateral temporal lobe, including posterior fusiform gyrus (Ranganath & Ritchey, 2012). PhC is more densely connected with posterior medial VVS regions, as well as scene sensitive regions such as retrosplenial cortex (Ranganath & Ritchey, 2012).

Another possibility is that these domains are present as organizing principles in MTL not simply because of downstream organization, but because these structures are components of broader neural systems involved in different functional aspects of perceiving and interacting with different categories or domains of objects. These systems may be reflected in long range functional and structural connectivity patterns with domain preferring or domain-organized cortex. Konkle et al. (2016) showed that the three prominent domain preference zones in the VVS (animate zone, large inanimate zone, and small inanimate zone) had different long-range connectivity patterns. For example, the dorsolateral zone that prefers small inanimate objects is more highly connected to parietal and motor regions involved in motor control. As a general framework, these regions may be part of functional systems that potentially are specialized for navigation, social interaction, and tool-use (for a recent review see Bi et al., 2016). Indeed, similar arguments have been made for MTL structures, based on their differential connectivity (Ranganath & Ritchey, 2012). PhC is connected to the medial parietal cortex, precuneus, ventrolateral parietal cortex and medial prefrontal cortex (default mode areas), while PrC is more densely connected with the amygdala, orbitofrontal cortex, and anterior ventrolateral temporal cortex (Kahn et al., 2008; Libby et al., 2012). This has led to the argument that these structures are part of systems that contribute differentially not just to memory processes, but to cognition overall. Specifically, according to this model PhC is part of a posterior-medial system (P-M system) essentially involved in creating situation models (i.e., a mental representation of the relationships between entities, actions, and outcomes) and PrC is a component of an anterior -temporal (A-T system) involved in assessing the significance of entities (people, objects). Therefore, perhaps animacy and real-world size organization in MTL structures more broadly reflects how these object representations may be functionally utilized, and how those structures contribute to those domains within a larger system. In support of this, recent work has shown that category-specific recognition memory signals are organized along a gradient by landmark suitability (Martin et al., 2018). Specifically, the gradient in MTL follows behavioral ratings, with objects that are the most suitable as landmarks, such as buildings and trees, represented more posteriorly in PhC, and planes more anterior (anterior portion of PhC and posterior portion of PrC). The authors suggest that landmark suitability is important to navigation, and that navigation could be the broader dimension of organization that shapes object recognition within PhC. Navigational relevance may be more broadly linked to the P-M system, in parallel with situation model building, however this relation remains to be further investigated.

An alternative, yet non-mutually exclusive possibility is that animacy and real-world size organization across MTL structures is driven by the perceptual and/or semantic features of the individual stimuli. This is a potentially interesting difference between MTL and more posterior VVS regions. Specifically, in more posterior VVS, category and domain organization may be functionally purposeful, for example, by enabling a fast readout of category-level information before identifying a distinct entity, which has adaptive value (see Grill-Spector & Weiner, 2014). In MTL structures, this organization may arise secondarily from individual object relations. A number of studies have shown that PrC is sensitive to distinctions among objects, even highly similar objects (Erez et al., 2012; Clark & Tyler, 2014, Tyler et al., 2013; Taylor et al., 2012; Kivisaari et al., 2012). Therefore, animacy and real-world size may reflect that distinct object responses vary in their feature space along these domains. In other words, distinct responses for animate objects overall may be more similar in feature space (and therefore neural representational space) than inanimate objects, with the same pattern for real-world size. Indeed, Clarke and Tyler (2014) showed that a model based on semantic features of individual objects, identical to the feature-norming approach described in Chapter 4, better fit the representational space in PrC in response to the same objects than a number of other models. These features were not broken down by perceptual or semantic attributes, however, making it difficult to determine whether perceptual, semantic, or a combination of feature types drives organization of object-based responses in PrC. Recent work by Martin et al. (*under review*) addressed this issue, as well as the influence of task on responses in both PhC and PrC. Specifically, they orthogonalized semantic and visual features for a set of object concepts, and scanned participants during either a semantic or visual-perceptual task while stimuli were presented visually as words. Interestingly, a combination of the visual and semantic model best fit the representational space in PrC, and there was no influence of task. By contrast, in PhC only the semantic model fit and only during the conceptual task. It would be interesting to understand how this pattern of findings aligns with the current pattern of results obtained for images of these objects across MTL structures. Is real-world size a more semantic dimension in general, and animacy a combination of visual and semantic dimensions? Evidence from studies with congenitally blind individuals speak to this idea. Both scene-preferring medial regions of the VVS, and small-object preferring dorsolateral occipital cortex are largely immune to visual experience and respond to many modalities of input, whereas selectivity for animate items in the lateral posterior fusiform gyrus requires visual experience, and is more closely tied to modality (for a review see Bi et al., 2016). A follow up study to examine MTL structure organization in congenitally blind and sighted individuals could be revealing. On the one hand, one might expect PrC organization to be dependent on visual experience, with less dependency for PhC and HpC. Alternatively, one might see that PrC is more dependent on visual experience than are PhC and HpC, but is less dependent on visual experience than is the lateral posterior fusiform, given that object representations in PrC are both visual and semantic in nature (as shown in Martin et al., *under review).*

5.3 Organization of recognition memory signals for objects in the MTL

Computational models of MTL function, such as those inspired by R-H theory, emphasize that memory signals are derived directly from representations (Cowell et al., 2006, 2010). Supporting this, Martin et al. (2013; 2017) showed category-specific recognition memory signals in PrC and PhC. These findings fit predictions from computational models in which a network is organized by stimulus features. Specifically, if stimuli that share features (i.e., are from the same category) share more overlap in neural representation, and recognition memory signals are a read-out of these representations, one would expect to pick up category-specificity in the memory signal. We extended this work, exploring category-specificity of recognition memory by taking advantage of our continuous recognition memory data. The benefit of our continuous task was that novel and familiar stimuli were identical, allowing us to circumvent the possibility that distinctions between novel and familiar stimuli are due to differences in stimuli sets, as in a traditional study-test paradigm in which novel items at test are necessarily a different stimuli set. We used this tightly controlled data set to ask whether we find evidence for the organization of memory signals by category and or domain for the large number of categories examined in our study. Further, we evaluated whether memory status was reflected in repetition suppression, a pattern-based change, or both. Surprisingly, we found no evidence of category or domain specificity in recognition-memory signals for objects. However, we did find robust repetition suppression in each MTL structure, as well as a pattern-based distinction between perceived novel and familiar responses across all object categories, that was not simply the result of repetition suppression.

It is somewhat surprising that we did not find any evidence of category or domain specificity in recognition memory signals for objects, given previous work in our lab. One possibility that needs to be ruled out in further analyses is that this lack of specificity can be accounted for by differences in power; notably, fewer trials inevitably went into the GLM when categories were split into novel versus familiar (maximum of 4 rather than 8 trials for each category). To address this issue, we can perform the analyses described in Chapter 2 on a sample of data that is based on 4 trials instead of 8, for example by using only novel or familiar responses, or a permutation of either. If we replicate our results from Chapter 2, that will support the idea that the lack of category and domain specific organization is not simply a power issue.

If power is not an issue and does not account for the results in Chapter 3, there are a number of alternative potential interpretations of the finding that memory-signals were not category or domain specific. One possible interpretation of these results is that it is simply not possible to assess distinctions between novel and familiar object representations when the stimuli are identical in both trial types. Another possibility is that in our task there was a contribution of recollection, or HpC feedback, when participants made their memory judgements. Indeed, if the original context of the novel encounter was reinstated during the memory decision, this might provide a shared memory signal across all object categories. This is possible in our study, because unlike in Martin et al.’s (2013; 2017) work in which they analyzed only familiarity responses, we did not employ a response procedure that distinguishes familiarity from recollection; instead participants simply responded with two options: new or old. However we did employ a response deadline procedure to encourage object familiarity over recollection, but as we did not explicitly rule out the contribution of recollection it remains a possible difference between the studies. Along these lines, it should be noted that we found across-object memory distinctions in the HpC, whereas Martin et al. did not. Although HpC involvement cannot be interpreted as recollection (a reverse inference), it is a possibility. Another alternative explanation is that a general memory signal is computed or reflected when the task space includes a large number of categories. The idea that different tasks either change neural computations, or change how fMRI reflects the underlying neural computations, has some evidence in support. For example, Davis et al. (2014) showed that representational space in anterior temporal lobe reflects category when analyzing stimulus evoked responses during a categorization task, but reflects individual differences in stimuli during a recognition-memory task (which requires item-based discrimination). Further, in the work described in the previous section by Martin et al. (*under review*). PhC showed task specific organization (i.e., conceptual organization in object evoked responses was observed during a conceptual task, but not during a perceptual task). It is possible that our task, which was associated with high performance due to the use of many different categories, did not emphasize category-specific memory processing the way previous work from our lab using a large number of exemplars from just a few categories did. Specifically, with a large number of different categories interleaved and less overall interference, perhaps distinct responses to individual stimuli were emphasized more than shared information across stimuli, leading to the across-category but not within-category recognition memory signal finding. There are a number of ways these ideas could be further tested. One could manipulate the amount of interference (i.e., the number of different categories, or similarity of objects within a category), or change the type or parameters of the memory task. In addition, one could use different feature selection techniques in these different task versions, such as selecting voxels that show the largest change (increases of decreases) between novel and familiar stimuli. Specifically, in a sharpening model of object repetition, a large number of units are activated more weakly at first presentation, and a smaller number of units more strongly upon repetition (Cowell et al., 2006). Following this logic, selecting voxels with the greatest absolute changes should narrow down the analysis to include voxels most involved in the representational change across novel and familiar presentations. If the network represents objects with overlapping features more similarly, one would expect that voxels involved in the sharpening of representations for objects from the same category would be more similar, leading to category-specific recognition memory signals. Presumably, voxels that show the greatest increase would also tap into these representations. Voxels with the greatest decrease might represent the voxels that drop out upon repetition, which may be informative in terms of a memory signal, but perhaps not in terms of stimulus specificity. Of course, caution should be taken when trying to relate fMRI data to computational models of sharpening, particularly when evidence of sharpening is usually at the neuronal level. For example, it is challenging to link increases or decreases of activity to excitatory or inhibitory neural activity, and a central premise of sharpening is that a set of neurons show increased firing for a representation with repetition, while the surrounding neurons are inhibitory (Kang et al., 2003). Overall, MVPA based approaches to analyzing object responses in fMRI data can aid in understanding how representations in MTL structures are organized, and how organization of representations is linked to their ability to contribute to object recognition memory.

5.4 Abnormal semantic knowledge in case of development amnesia

In Chapters 2 and 3, we investigated differences in content across MTL regions, and how this content contributes to recognition memory for objects. When thinking about content in the HpC, both R-H theory and the BIC model propose that HpC contains multidimensional representations that include objects and their spatio-temporal context. While the BIC model emphasizes the contribution that binding of items-in-context makes to episodic memory, focusing on how these types of representations could contribute processes beyond episodic memory led us to the novel idea that the HpC might be involved in object concepts, an important aspect of semantic memory. It has long been thought that the HpC contributes to episodic memory, but not semantic memory. Influential evidence for this theory came from the finding that individuals with developmental amnesia, resulting from relatively selective HpC damage occurring early in life, had normal semantic knowledge despite their severe impairment in episodic memory. In Chapter 4, we challenged this sharp distinction by testing the hypothesis that the HpC is involved in some aspects of object knowledge not previously considered or tested. Specifically, we hypothesized that an item-in-context binding mechanism that serves episodic memory is also important for acquiring extrinsic feature-knowledge of objects, features that are particularly important for objects in the nonliving domain. We found that HC, an individual with developmental amnesia, produced fewer extrinsic features, despite producing a normal number of intrinsic features. NB, an individual with damage to the PrC and EC, but not HpC, by comparison did not differ from controls on intrinsic or extrinsic feature generation. Further, HC had abnormal typicality ratings specific to nonliving objects, objects that rely more heavily on extrinsic features. NB on the other hand, did not differ from controls in her typicality ratings for living or nonliving objects.

One key aspect of the interpretation of this study is that the HpC is necessary for gaining extrinsic feature knowledge for objects throughout development, but that conceptual knowledge becomes independent of the HpC as it becomes a part of the object concept itself. This view is supported by the complementary learning systems model of HpC and cortex, which emphasizes that information from episodes is transferred by the HpC to cortex, and often becomes schematized, or more semantic in nature (Norman &O’Reilly, 2003). While the CLS view does not specify the part of cortex to which information is transferred, converging evidence suggests that both anterior medial temporal cortex (i.e., PrC and PhC) and lateral anterior temporal cortex may contain object concept representations. In general, the anterior temporal lobes are the target for damage in semantic dementia (Hodges et al., 1992; Patterson et al., 2007), and transcranial magnetic stimulation in healthy adults slows down naming of living and nonliving objects (Pobric et al., 2007). Further, some evidence suggests that semantic knowledge of extrinsic features for nonliving objects can be decoded from the anterior temporal lobes (Peelen and Caramazza, 2012). In terms of PrC, both visual and semantic object information is represented there (Martin et al., *under review*; Clark & Tyler, 2014; Bruffaerts et al., 2013), whereas in PhC object concept information is represented in the context of a conceptual task (Martin et al., *under review*). In Chapter 2 and 3 of the current dissertation work, while we did not focus on semantic representations, we do find that object representations for a set of nonliving object categories in PhC, and object evoked representations in both PhC and PrC are organized by whether an object is animate or inanimate. This converging evidence suggests that both anterior medial temporal cortex (i.e., PrC and PhC) and lateral anterior temporal cortex may contain object concept representations. However, the relationship between these representations and HpC function in potentially shaping the development of these representations is still an open area of inquiry.

One way to further test ideas about HpC contributions to object concepts would be to evaluate whether young children, with protracted HpC development, have less extrinsic feature knowledge early on, which increases over time in relation to their HpC development or to connectivity between HpC and cortex. More specifically, one could test whether their object evoked representational space in PrC, PhC, or lateral anterior temporal cortex was organized by intrinsic features (a model of object relation based on intrinsic features only) early on, while later in development was best modeled by a combination of extrinsic and intrinsic features. Further, one could see whether these changes in representational space correlate with HpC development or connectivity with cortex.

While this dependence on the HpC for acquiring extrinsic object knowledge might be most evident in development, when a wealth of new semantic knowledge is gained, it is possible the HpC continues to be needed for acquiring new extrinsic knowledge. For example, work by Duff and colleagues has shown decreased learning of new vocabulary in patients with HpC damage. Specifically, individuals with HpC damage had less vocabulary depth and semantic richness, both measures that track increasing semantic knowledge after initial acquisition of a concept (Klooster et al., 2013). One can consider linking an orthographic word to new meaning as a similar process to linking an object to an extrinsic feature (Konkel et al., 2008). These results would then suggest that even after a basic object concept is established (i.e., an object and some key concrete features), further elaboration of that concept (i.e., adding additional extrinsic or abstract features) might continue to involve the HpC over the lifetime. In our study, we focused on concrete object concepts and features that come to mind immediately, which we presume are learned fairly early during development. However, an interesting prediction to test is whether more unusual or abstract features of object concepts require the HpC even in adulthood.

A final open question is whether the HpC is needed to access previously learned extrinsic features. For example, it is not known whether the HpC needs to come online for such access in feature production tasks or typicality judgements. A nice follow up study to address this question would be to test adult-onset amnesics on our feature generation and typicality task. Specifically, we could test the hypothesis that adult-onset amnesics would perform normally on feature generation and typicality tasks, because concrete extrinsic features were acquired during development when they had an intact HpC, and this object-knowledge was transferred to cortical representations that can still be accessed. This dissociation between developmental and adult onset HpC damage would strengthen our interpretation that the HpC is involved in the acquisition of concrete feature knowledge, and provide further support for a CLS model of HpC function (Norman & O’Reilly, 2003).

5.5 Benefits of exploring representational space and future directions

The current findings add to a growing body of literature employing representational similarity analysis as an approach. Evaluating representational space is a fruitful approach for exploring the content of different brain structures, in order to better understand how structures contribute to broad aspects of cognition and behavior. One way to move forward with this approach is to push tracking changes in representational space. For example, task based influences on the space can help us to understand the flexibility or rigidity of a brain structures functional architecture. Further, examining changes in representational space across development and how these changes relate to structural and functional connectivity may provide insight into key differences between different structures that underlie their shared or differential contributions to aspects of cognition and behavior. Lastly, evaluating stimulus-evoked space across structures in unique neurological conditions, such as in individuals who are congenitally blind, or in individuals with known abnormalities in semantic knowledge, such as in semantic dementia, can help both to inform our basic understand of neural function, and to enhance our understanding of these unique conditions.

5.6 Conclusions and future directions

The overarching goal of this dissertation research was to increase our understanding of the functional specialization and organization of MTL structures. Our approach was to move away from dichotomies between processes, such as memory versus perception, familiarity versus recollection, or episodic versus semantic memory, and instead focus on content and how content contributes to different aspects of cognition and behavior. We chose to focus on objects, and investigating content in relation to object category or domain, given this has led to a rich understanding of neural organization in the VVS. Further, while we know object perception and memory rely crucially upon MTL structures, a detailed model of how they contribute is still underway. We were guided by the R-H view, and evidence that a crucial distinction between brain areas may be in the representational complexity of content, with more anterior areas containing representations of higher dimensionality. However, the relation between the dimensionality of representations and differences in category or domain sensitivity and organization has yet to be well linked. Our attempts to map out the representational space in each structure during stimulus presentation, and in recognition memory signals, is an important step in this direction.

Our findings on MTL content organization add further refinement to the current theories discussed, R-H theory and the BIC model. Similar to previous work in our lab, we show that the BIC model of PhC function is not fully comprehensive, as PhC is involved in object processing as opposed to only spatial episodic context. Instead, PhC may be better characterized by the way it contributes to both object and context representations.

Understanding how content in PhC changes across experiences, such as in recollection versus semantic knowledge use versus navigation, may help to create a cohesive model of this structure. The same is true for PrC, a structure that clearly contains high-level object representations, and as we show is organized by animacy but not real-world size. Further understanding the limitations of PrC object sensitivity or generality, as well as how these representations contribute to object perception versus recollection will help to get a full picture of this structure. Overall, a useful approach moving forward, is to understand content differences in PhC and PrC, by combining system-based and feature-based models. Beyond this, an exciting and important aspect of understanding the functions of these structures will be to understand how their content is utilized for different processes.

This approach of using content as a way to characterize and understand the function of a structure becomes somewhat more challenging when examining the HpC. Taking an R-H view of the HpC, and focusing on how HpC content contributes to different aspects of cognition and behavior may provide an interesting challenge to both the R-H and the BIC model. For R-H view, it is a potential challenge as to whether we should think of the representational content (item-in-context) as a hippocampal representation per se. For the BIC model, the challenge concerns expanding the scope of binding items-in-context beyond a contribution to episodic memory. The R-H theory view of HpC is similar to the BIC view, in that it emphasizes HpC representations as associative, and often highly dimensional combinations of multiple objects and their spatio-temporal context (Cowell et al., 2010). However there is a difference in emphasis between the models, with R-H view emphasizing the type of representation, and the BIC model emphasizing *binding* of these elements as a key function. This brings us to a complicated question, as to whether HpC contains representations that are a mix of object and context information, or whether the representation in the HpC itself is abstracted away from its input, but serves to bind or link that input in cortex. Along these lines, theories have emphasized binding as a key computation that allows for linking cortical elements into an episodic representation, rich in detail and experienced holistically, despite the fact that item and context information within the episode are cortically distributed. Initially binding these elements during encoding of an episode can lead to later reactivation, thought to be the neural basis of recollection. However it is unclear whether HpC contains a rich multidimensional representation of the episode itself, or creates an orthogonalized tag (sometimes described as a “bar code” or a pointer) for the distinct experience. Either one of these possibilities could explain the general finding that the HpC is agnostic to category or domain level organization of representations. If the HpC contains item-in-context representations, they likely combine a mix of different stimuli information which would not allow for easily capturing shared information across experiences, that is, category specific information. In the latter case, the “tags” for different episodes are thought to be orthogonalized to detract from interference, which would also decrease any shared stimuli information across representations. Importantly, in Chapter 2, we show that representations in the HpC are not entirely agnostic, with a trend for buildings and a broader organization by whether objects are large or small, and other studies have found the HpC is sensitive to scene stimuli. This suggests that at least for some stimuli in some circumstances HpC representations do reflect shared stimuli information, showing that further research into HpC content is needed.

In general, focusing on content (either contained in, or orchestrated by) of the HpC as opposed to processes, led us to the novel hypothesis that binding of items- in-context is important for semantic memory. We found evidence confirming this hypothesis, presented in Chapter 4. This finding suggests that there is benefit in moving beyond and episodic-memory only focused view of the HpC, such as in the BIC mode. HpC function is not fully captured by considering representations within HpC during a task, as one might following an R-H theory approach either - as we show an important contribution of HpC to the development of representations presumably through interactions with cortex and binding. In other words, whether we consider binding items-in-context an HpC representation *or* a computation, it is important to consider the different aspects of learning and memory this structure can contribute to beyond one process (i.e., episodic memory). Our finding that the HpC may be especially important for concept knowledge in the nonliving domain is at the level of *forming* associative representations. It is important to examine this in development, where we can better assess whether representations within the HpC are contributing to object concepts, or representations are created by interactions between HpC and cortex. Specifically, one could ask whether initially these representations are created and maintained in the HpC and then transferred to cortex? Or are these representations always created through HpC cortical interactions that allow for extrinsic feature information to be linked to objects, and ultimately integrated intro cortical object concept representations?

Considering whether the HpC “contains” or “creates” multidimensional representations, and how this might change at different stages of development, brings up an important consideration when characterizing the function of MTL regions by the R-H or BIC model. Neither of these models capture the highly dynamic and potentially distributed basis of representations, or the different way representations are used when say, perceiving the world, or recollecting an episode. One way to address this within the realm of the current methods is to explore how representational space changes across different tasks, and, potentially to use more naturalistic stimuli to evoke representational space. Another important aspect to keep in mind is that, while much evidence suggests a hierarchy in representational complexity along the VVS, this information pathway is full of back projections and parallel projections (Kravitz et al., 2012). Further, a wealth of fascinating evidence suggests that top down or feedback effects can shape representations in earlier areas (Hindy et al., 2016; Rao & Ballard, 1999). Using techniques with high temporal resolution in parallel with fMRI techniques to investigate representational changes in a feedforward, feedback, and parallel manner is therefore highly important. Additionally, while focusing on the representations and/or computations in individual brain structures is important, it is also crucial to situate the patterns of results within structures to their corresponding larger-scale networks, in order to fully understand their contributions to cognition and behavior. A good example of this is the anterior-temporal and posterior-medial model of MTL structures (Ranganath & Ritchey, 2012). One can scale up even further, by looking at how these networks interact. This is especially important, given that our experiences and memories often involve animate objects, small and large inanimate objects, and a rich spatial and temporal context that we experience as an integrated and dynamic whole, for example when we imagine or remember meeting our friends in the park for a picnic.

Appendices

Appendix B: Documentation of ethics approval for Chapter 2 and 3

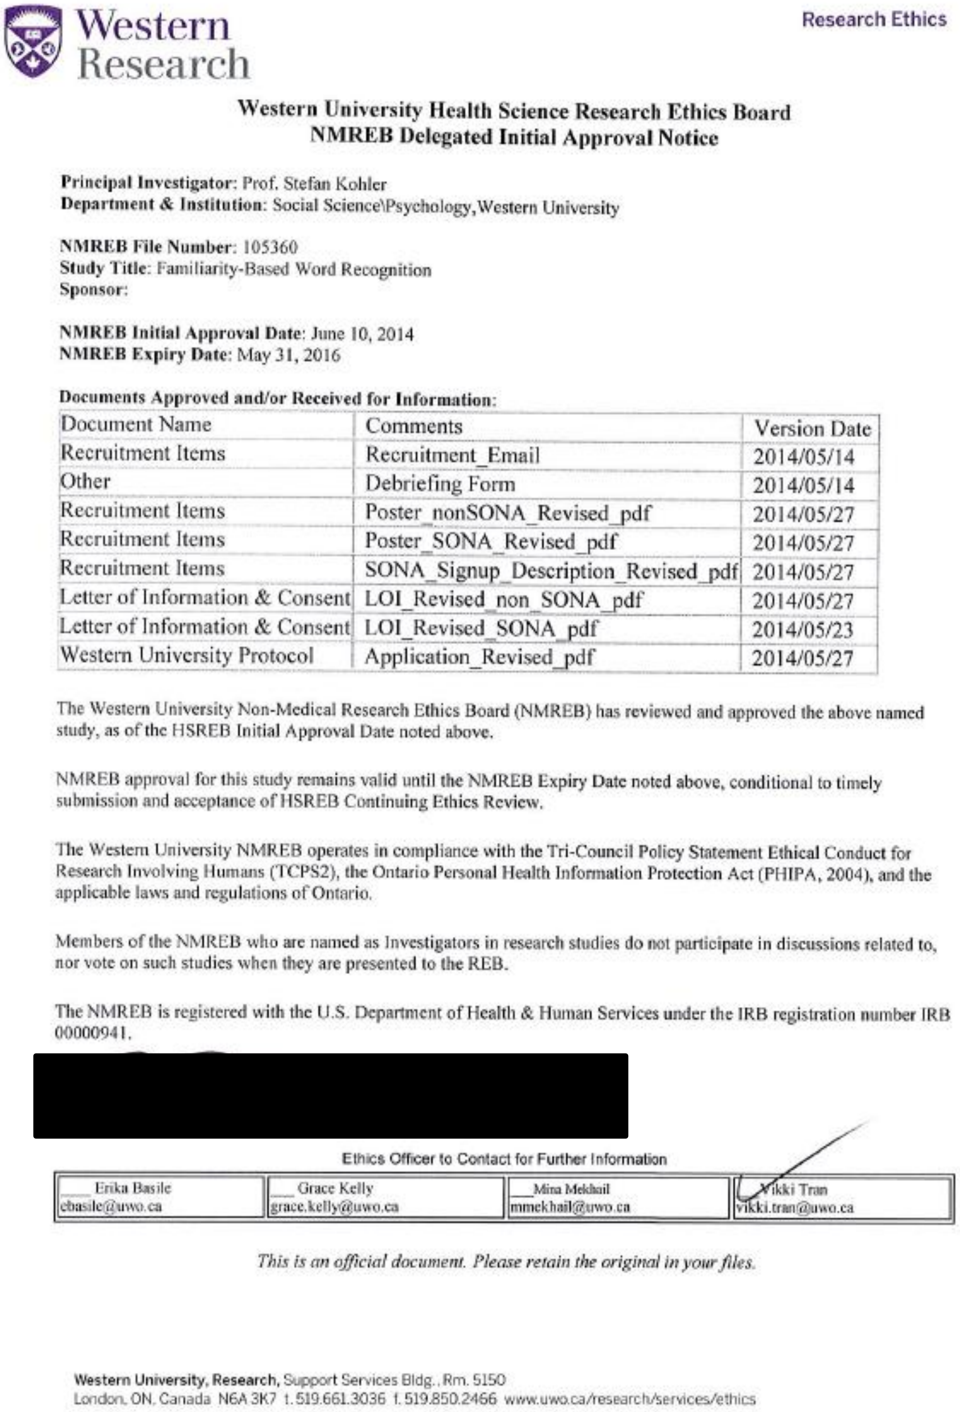

Western Research

Research Ethics

Western University Health Science Research Ethics Board
NMREB Delegated Initial Approval Notice

Principal Investigator: Prof. Stefan Kohler
Department & Institution: Social Science\Psychology,Western University

NMREB File Number: 105360
Study Title: Familiarity-Based Word Recognition
Sponsor:

NMREB Initial Approval Date: June 10, 2014
NMREB Expiry Date: May 31, 2016

Documents Approved and/or Received for Information:

Document Name	Comments	Version Date
Recruitment Items	Recruitment_Email	2014/05/14
Other	Debriefing Form	2014/05/14
Recruitment Items	Poster_nonSONA_Revised_pdf	2014/05/27
Recruitment Items	Poster_SONA_Revised_pdf	2014/05/27
Recruitment Items	SONA_Signup_Description_Revised_pdf	2014/05/27
Letter of Information & Consent	LOI_Revised_non_SONA_pdf	2014/05/27
Letter of Information & Consent	LOI_Revised_SONA_pdf	2014/05/23
Western University Protocol	Application_Revised_pdf	2014/05/27

The Western University Non-Medical Research Ethics Board (NMREB) has reviewed and approved the above named study, as of the HSREB Initial Approval Date noted above.

NMREB approval for this study remains valid until the NMREB Expiry Date noted above, conditional to timely submission and acceptance of HSREB Continuing Ethics Review.

The Western University NMREB operates in compliance with the Tri-Council Policy Statement Ethical Conduct for Research Involving Humans (TCPS2), the Ontario Personal Health Information Protection Act (PHIPA, 2004), and the applicable laws and regulations of Ontario.

Members of the NMREB who are named as Investigators in research studies do not participate in discussions related to, nor vote on such studies when they are presented to the REB.

The NMREB is registered with the U.S. Department of Health & Human Services under the IRB registration number IRB 00000941.

[redacted]

Ethics Officer to Contact for Further Information

___ Erika Basile ebasile@uwo.ca	___ Grace Kelly grace.kelly@uwo.ca	___ Mina Mekhail mmekhail@uwo.ca	✓ Vikki Tran vikki.tran@uwo.ca

This is an official document. Please retain the original in your files.

Western University, Research, Support Services Bldg., Rm. 5150
London, ON, Canada N6A 3K7 t. 519.661.3036 f. 519.850.2466 www.uwo.ca/research/services/ethics

Appendix B: Documentation of ethics approval for Chapter 4

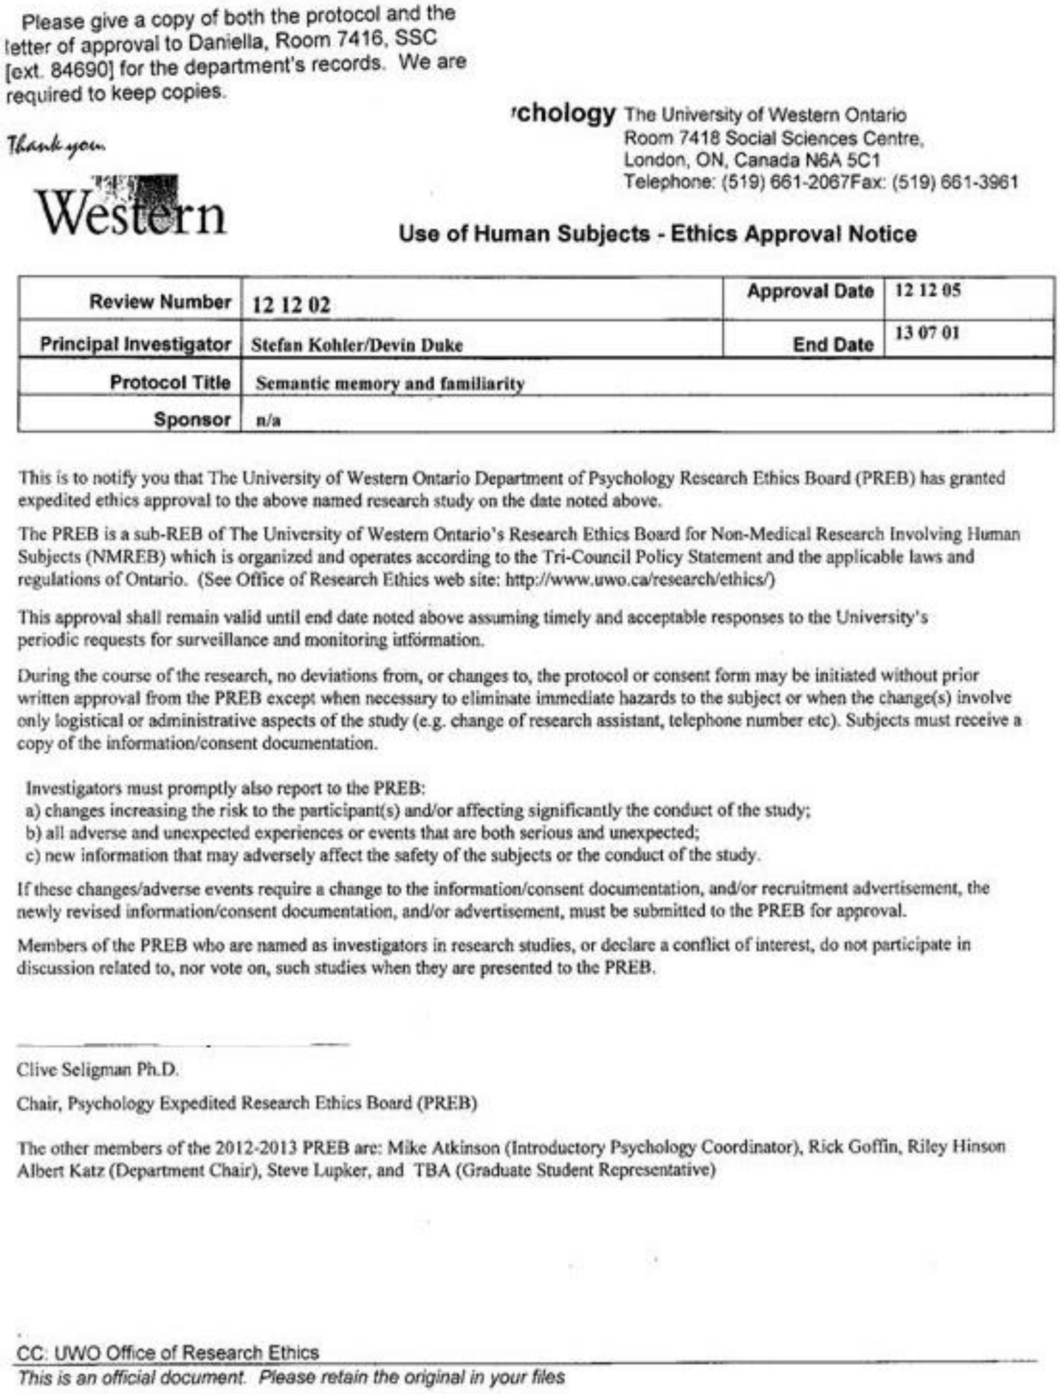

Please give a copy of both the protocol and the letter of approval to Daniella, Room 7416, SSC [ext. 84690] for the department's records. We are required to keep copies.

Thank you.

Western

chology The University of Western Ontario
Room 7418 Social Sciences Centre,
London, ON, Canada N6A 5C1
Telephone: (519) 661-2067Fax: (519) 661-3961

Use of Human Subjects - Ethics Approval Notice

Review Number	12 12 02	**Approval Date**	12 12 05
Principal Investigator	Stefan Kohler/Devin Duke	**End Date**	13 07 01
Protocol Title	Semantic memory and familiarity		
Sponsor	n/a		

This is to notify you that The University of Western Ontario Department of Psychology Research Ethics Board (PREB) has granted expedited ethics approval to the above named research study on the date noted above.

The PREB is a sub-REB of The University of Western Ontario's Research Ethics Board for Non-Medical Research Involving Human Subjects (NMREB) which is organized and operates according to the Tri-Council Policy Statement and the applicable laws and regulations of Ontario. (See Office of Research Ethics web site: http://www.uwo.ca/research/ethics/)

This approval shall remain valid until end date noted above assuming timely and acceptable responses to the University's periodic requests for surveillance and monitoring information.

During the course of the research, no deviations from, or changes to, the protocol or consent form may be initiated without prior written approval from the PREB except when necessary to eliminate immediate hazards to the subject or when the change(s) involve only logistical or administrative aspects of the study (e.g. change of research assistant, telephone number etc). Subjects must receive a copy of the information/consent documentation.

Investigators must promptly also report to the PREB:
a) changes increasing the risk to the participant(s) and/or affecting significantly the conduct of the study;
b) all adverse and unexpected experiences or events that are both serious and unexpected;
c) new information that may adversely affect the safety of the subjects or the conduct of the study.

If these changes/adverse events require a change to the information/consent documentation, and/or recruitment advertisement, the newly revised information/consent documentation, and/or advertisement, must be submitted to the PREB for approval.

Members of the PREB who are named as investigators in research studies, or declare a conflict of interest, do not participate in discussion related to, nor vote on, such studies when they are presented to the PREB.

Clive Seligman Ph.D.

Chair, Psychology Expedited Research Ethics Board (PREB)

The other members of the 2012-2013 PREB are: Mike Atkinson (Introductory Psychology Coordinator), Rick Goffin, Riley Hinson Albert Katz (Department Chair), Steve Lupker, and TBA (Graduate Student Representative)

CC: UWO Office of Research Ethics

This is an official document. Please retain the original in your files

Curriculum Vitae

EDUCATION

Ph.D., Psychology (in progress)

University of Western Ontario

Brain and Mind Institute (BMI)

Track: Behavioral and Cognitive Neuroscience

Laboratory: Memory Lab in Cognitive Neuroscience

Principal Investigator: Dr. Stefan Köhler

M.S., Cognitive Science, December, 2013

University of Trento/SISSA (International School for Advanced Studies, Trieste)

Double Degree Fellowship

Track: Cognitive Neuroscience

Laboratory: Brain and Language

Principal Investigator: Dr. Giosuè Baggio

B.A., Psychology, May, 2010

Drew University, Madison, NJ

Minors: Classical Studies & Philosophy

G.P.A.: 3.7

HONORS & AWARDS

University of Trento Merit Award: 2000 euro

Ontario Trillium Scholarship for Doctoral Study in Canada: 160,000$

Erasmus Lifelong Learning Programme Scholarship: 1500 euro

UNITN/SISSA Fellowship Award: 18,000 euro

Faculty for Undergraduate Neuroscience and Nu Rho Psi Travel Award: 750$

James A. McClintock Award for Excellence in Psychology: 700$

Magna Cum Laude

Psy Chi (International Honor Society in Psychology)

Nu Rho Psi (National Honor Society for Neuroscience)

ADDITIONAL RESEARCH EXPERIENCE

BCCN (Bernstein Center for Computational Neuroscience), Erasmus Internship (4/13-7/13)

Nathan Kline Institute for Psychiatric Research, Research Assistant (7/10-9/11)

University of Minnesota REU Program in Behavioral and Cognitive Sciences (6/09-8/09

Drew University Behavioral Neuroscience Laboratory, Directed Research (1/08-5/08)& Research Assistant (10/07-11/08)

Drew Summer Science Institute (DSSI) (7/08-9/08)

PUBLICATIONS

Blumenthal, A., Stojanoski, B., Martin, C., Cusack, R. & Köhler, S. Animacy and real-world size shape object representations in the medial temporal lobes. (*under review, Human Brain Mapping*).

Blumenthal, A., Duke, D., Bowles, B., Gilboa, A., Rosenbaum, S., Köhler, S. & McRae, K. Abnormal semantic knowledge in a case of developmental amnesia. *Neuropsychologia* (*in press*).

*Linking Words and Memory: How we Remember the Structure of Things. Story for

CNS website. Written by: Lisa Munoz.

https://www.cogneurosociety.org/linking-words-and-memories-how-we-remember-the-structure-of-things/

Mostafa Kia M., Pedregrosa, F., **Blumenthal, A.** & Passerini, A. Group-level spatio-temporal pattern recovery in MEG decoding using multi-task joint feature learning. *Journal of Neuroscience Methods*, 285, 97-108.

Baggio, G., Cherubini, P., Pischedda, D., **Blumenthal, A**., Haynes, JD. & Reverberi, C. Multiple neural representations of elementary logical connectives. *Neuroimage,* 135, 300-310, 2016.

Blumenthal, A., Steiner, A., Seeland, K. & Redish, A. D. R. Effects of pharmacological manipulations of NMDA-receptors on deliberation in the multiple T-task. *Neurobiology of Learning and Memory*, 95, 376-384, 2011.

Cousens, GA, Skrobacz, CG & **Blumenthal, A**. Nucleus accumbens carbachol disrupts olfactory and contextual fear-potentiated startle and attenuates baseline startle reactivity. *Behavioural Brain Research,* 216, 673-680, 2010.

INIVITED TALKS

CSBBCS (Canadian Society for Brain, Behaviour and Cognitive Science) 2016. Symposia, Categorization: Causes and Consequences. *Object organization in the medial temporal lobe.*

TAMeG Spring Meeting 2016. Toronto Area Memory Group. *Object organization in the medial temporal lobe.*

University of Waterloo 2016. *Object organization in the medial temporal lobe.*

SRPoiSE 2016: The Consortium for Socially Relevant Philosophy of/in Science and Engineering. *The Lab Associates Program: Graduate Training Towards Interdisciplinary work with Neuroscientists*. Panel: Jessey Wright, Robert Foley, Rob Whitwell, **Anna Blumenthal**. Talk: Philosophical Impact on Neuroscience

SELECTED CONFERENCE PRESENTATIONS

Blumenthal, A., Stojanoski, B., Martin, C., Cusack, R. & Köhler, S. Organization of object representations in the medial temporal lobe. CNS (Cognitive Neuroscience Society) Annual Meeting, San Francisco, CA, 2017.

Blumenthal, A., Duke, D., Bowles, B., Gilboa, A., Rosenbaum, S., Köhler, S. & McRae, K. Abnormal semantic knowledge in a case of developmental amnesia. LOVE (Lake Ontario Visionary Establishment) Meeting, Niagara Falls, ON. 2017.

Blumenthal, A., Stojanoski, B., Martin, C., Cusack, R. & *Köhler, S.* Representational similarity analysis of category-related recognition-memory signals in the human medial temporal lobe. CAN-ACN 2016 (Canadian Association of Neuroscience) Meeting, Toronto, ON. 2016

Blumenthal, A., Stojanoski, B., Martin, C., Cusack, R. & *Köhler, S.* Representational similarity analysis of category-related recognition-memory signals in the human medial temporal lobe. Vision Sciences Society Sixteenth Annual Meeting, St. Pete Beach, FA. 2016

Blumenthal, A., Steiner, A., Seeland, K. & Redish, A. D. R. Involvement of NMDA-receptors in vicarious-trial-and-error behaviors on a spatial task. Society for Neuroscience Annual Conference. 2010.

Blumenthal, A., Boska C. R., Greger-Moser M., Hull-Rawson L.H. Miele, F. Modulation of acoustic startle by appetitive and aversive conditioned olfactory cues. Society for Neuroscience Annual Conference. 2008.