NORTHWESTERN UNIVERSITY

Memory Specificity and Generalization through Neural Replay during Sleep

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

For the degree

DOCTOR OF PHILOSOPHY

Field of Psychology

By

Sarah Witkowski

EVANSTON, ILLINOIS

September 2020

ABSTRACT

Memory Specificity and Generalization through Neural Replay during Sleep

Research on how sleep contributes to memory has blossomed in recent years. These studies have generally focused on whether or not sleep impacts various types of memory independently. An open question is whether sleep interactively influences different memory types. My research focuses on two types of memory—specificity and generalization. Whereas we forget many memories encoded during our day, sleep can help counteract some of that forgetting and protect specific memories. However, we don't know whether sleep improves generalization at the cost of specificity or vice versa. Based on prior work, I hypothesized that specificity and generalization exhibit a trade-off such that remembering detailed specifics blocks generalization. Additionally, I hypothesized that when generalization occurs, the specifics of memories are lost.

In this dissertation, I describe two experiments designed to address this question. Chapter three is composed of two studies where I showed participants videos of outdoor scenes that fall into three categories. After learning, participants were tested on recognizing old or new videos (specificity) of the terrains as well as identifying to which of three highly-confusable categories the video belonged (generalization). In the first study of chapter three, participants had either a 12-hour delay that included sleep or was comprised of wake only. In the follow-up study we controlled circadian factors by having participants remain awake or nap for 90 minutes during the afternoon while we recorded brain activity. After the delay, they were once again tested on specificity and generalization. We did not find a significant difference between the sleep and wake groups in either of the experiments, but did find that forgetting of specific videos happened more rapidly than forgetting of categorization knowledge. Although these results did not provide

evidence for the hypothesized specificity-generalization tradeoff, they did provide support for the prototype theory of memory as opposed to exemplar theory.

In chapter four, I focused on a period of afternoon sleep after participants attempted to memorize a set of paintings created by six different artists. I used targeted memory reactivation (TMR) with sounds associated with each artist in an attempt to modulate memory processing during sleep. The goal was to test whether memory reactivation alters trade-offs between memory specificity and generalization. I hypothesized that TMR cues would improve memory for generalization over remembering specific paintings. While cuing did not alter participants' performance on generalization, it did impact memory for specific, studied paintings. Participants were worse at identifying cued paintings compared to uncued paintings after the nap, possibly because memory reactivation during sleep facilitated storage of generic aspects of paintings at the expense of some details.

These experiments provide evidence for several ways in which different neural mechanisms operative during sleep support memory for specifics and generalization. The results show that memory for specifics is forgotten at a faster rate than generalization, and that directed replay during sleep can impact that rate. These findings should guide future research, particularly with regard to developing experimental designs that can independently assess the two memory types. The experiments shed light on the varying forgetting rates across different memory types and provide a springboard for future research into mechanisms that govern memory consolidation during sleep.

ACKNOWLEDGEMENTS

Completing this PhD has been a labor of love, supported by a wide network of family, friends, and colleagues. It’s overwhelming how many people encouraged me on both my graduate journey and through these experiments. I want to acknowledge the people who made the biggest impact on my voyage by including them here.

First of all, I want to thank my advisor Dr. Ken Paller for his guidance throughout my training. Ken has spent hours guiding me through experimental design planning, data analysis, and manuscript editing. But he’s gone much beyond that, teaching me how to juggle, introducing me to fellow scientists & opportunities, and allowing me the freedom to explore my interests- academic and science writing.

I would also like to thank my committee, Dr. Marcia Grabowecky, Dr. Paul Reber, and Dr. Ali Preston, for their guidance and advice on the projects included in this document. From the inception of my projects, I’ve been incredibly thankful to my committee’s feedback and assistance. Additionally, thank you to the members of the Cognitive Neuroscience Lab: Jessica Creery, Laura Batterink, Eitan Schechtman, Susan Florczak, Bruce Caughran, Nathan Whitmore, Adrianna Bassard, and Karen Konkoly for creating a science-loving lab where I’ve felt perfectly at home.

I also want to recognize and thank my research assistants. Thank you to Leonardo Dionisio for putting up with my scatter-brained meetings as my first research assistant and teaching me how to be a mentor. Thanks also to Janessa Engelmeyer and Victoria Lee for all the work they put into helping develop and run experiments with me.

Outside of my academic world, I've also had an incredible support network seeing me though on this journey. Firstly, my family deserve recognition for all their support and cheering me on through the hard times. My father, Doug Witkowski, taught me to love science while my mom, Sally Witkowski, taught me to press forward in the hard times because there's always light at the end of the tunnel. My brother, Douglas Witkowski, has been extremely patient putting up with my long neuroscience rants over phone calls and holidays and deserves a medal for it.

I also want to acknowledge the support and love of my fiancé, Daniel Regueira. When I decided to move up to Chicago for a PhD, he was there to help me pack and follow me to the Midwest. And of course, a massive thank you to our pup Zoey who has seen me though many a long writing session and offered the best snuggles right when I needed them.

To Ben Reuveni, who I met during interview weekend without realizing I'd found my closest friend, thank you so much for keeping me fit. Our thrice weekly lifting sessions kept me in shape and provided the perfect time to discuss our research in-depth. I can't begin to tell you how many times I returned to the lab, newly inspired to try a new analysis or framing of a chapter.

Finally, I want to thank three groups of friends for their support for the past six years. Moving to Chicago was a big step away from my home in Texas, but my cohort immediately made me feel right at home. We studied together, threw ridiculous Halloween parties, and provided emotional support during academic struggles. I could not have asked for a better group of psychology researchers than the ones in my 2014-2020 (or 2019) cohort. The next group I want to thank is my trivia crew. For six years, we've gone to Evanston Pub to eat, drink, and merrily try to out-do each other with our trivial knowledge. Having a Wednesday to look forward

to with fellow psychology graduate students has kept me motivated throughout many a long week. Thank you. And lastly, thanks to Erin Anderson, Francisco Maravilla, Lorien Elleman, Nathan Couch, Hollen Reischer, Natalie Gallagher, and Cassie Brandes for being the best dang podcast commune a girl could ask for. Without PCC, I would have missed out on so many great conversations, engaging podcasts, and tasty meals. But most importantly, the community provided by our group powered me through the craziness of graduate school.

I know there are so many other people who have impacted my journey that I could fill a document nearly as long as this dissertation. I feel so grateful for everyone that has changed my life and seen me through this experience. Thank you.

Table of Contents

ABSTRACT 3

ACKNOWLEDGEMENTS 5

CHAPTER 1 11

Memory Research 11

Memory Consolidation and Sleep 15

Endogenous memory replay during sleep 18

Targeted Memory Reactivation 25

Completed Studies 29

CHAPTER 2 30

Introduction 30

Sleep's impact on declarative memory for specifics 33

Memory generalization during sleep 34

Contrasting specificity and generalization during sleep 36

Future prospects using targeted memory reactivation 38

CHAPTER 3 42

METHOD 45

STUDY 1 RESULTS AND DISCUSSION 53

STUDY 2 RESULTS AND DISCUSSION 61

GENERAL DISCUSSION 69

CHAPTER 4 74

METHOD 80

RESULTS 85

DISCUSSION 91

CHAPTER 5 96

Summary 96

Future Directions 98

Conclusion 101

REFERENCES 103

LIST OF TABLES AND FIGURES

Table 3.1. Sleep group's physiology and minutes in sleep stages.

Table 4.1. Minutes spent in sleep stages.

Figure 2.1. Schema of memory consolidation for memory specificity and generalization.

Figure 3.1. Examples of alpha, bravo, and charlie videos.

Figure 3.2. Representation of categorical stimuli in 3D space.

Figure 3.3. Graph of training blocks in study 1.

Figure 3.4. Graphs of recognition accuracy and confidence in study 1.

Figure 3.5. Graphs of categorization accuracy and confidence in study 1.

Figure 3.6. Graph of training blocks in study 2.

Figure 3.7. Graphs of recognition accuracy and confidence in study 2.

Figure 3.8. Graphs of categorization accuracy and confidence in study 2.

Figure 4.1. Flowchart of experimental steps.

Figure 4.2. Examples of learning and generalization and specificity tests.

Figure 4.3. Graph of generalization performance before and after sleep.

Figure 4.4. Graph of memory specificity performance before and after sleep.

Figure 4.5. Graph of reaction times for generalization and memory specificity.

Figure 4.6. Graphs of memory performance across three time points.

Figure 5.1. Examples for testing boundary extension.

CHAPTER 1
General Introduction

Memory Research

Unlike the hard drive of a computer, memory isn’t a simple filing system that precisely records everything as soon as it happens. How memories are integrated and stored in a retrievable fashion is an important question in memory research. Across the years, researchers have proposed several theories to explain how memories are stored, integrated, and retrieved. Memory consolidation refers to the process of stabilizing and integrating memories within the brain, although there is still disagreement on how this is achieved.

Standard Memory Consolidation Theory

In the case of declarative memory, one proposed mechanism of memory consolidation relies specifically on interactions between the hippocampus and neocortex (O’Reilly & Norman, 2002). Memories initially encoded with the help of the hippocampus are replayed and communicated to the neocortex in order to create more permanent, stabilized forms of the memory. This theory is generally referred to as the standard memory consolidation theory. Importantly, this model implies that more recent memories would still retain neural connections to the hippocampus although this connection is not needed for retrieval after enough time has elapsed (Dagmar Zeithamova et al., 2012).

More recent studies have expanded on this idea, demonstrating that the anterior and posterior regions of the hippocampus itself might relate to whether information stored within the hippocampus is highly detailed or more generalized (Moscovitch et al., 2016). The anterior hippocampus is more connected with the ventromedial prefrontal cortex and is associated more with pattern completion while the posterior portion of the hippocampus has more connections to the parietal lobe and is more associated with pattern separation. Regardless of the hippocampal region involved, the communication of information between the hippocampus and the neocortex is crucial for memory consolidation. However, another major competing theory also takes this hippocampal data into consideration.

Multiple Trace Theory

Rather than conceiving of the hippocampus as a temporary component that is needed to reference cortical memories before they are stored long-term in the neocortex, some researchers fall into a different camp. Multiple trace theorists view the brain as akin to a wildlife path found in the woods. If a deer walks along a particular route once, there's very little impact made on the grass and earth. But if that same deer (or multiple deer) continue to use that route, the grass will give way to dirt and the path will become more defined. In neural terms, multiple trace theory is a system where the hippocampus is always involved in memory recall (Nadel et al., 2000). Every time a memory is encoded, a memory trace between the hippocampus and the neocortex is laid down (Moscovitch et al., 2005). In the case of multiple overlapping memories, a gist representation of the information is stored in the neocortex and doesn't necessarily need the hippocampus. In this respect, both multiple trace theory and standard memory consolidation theory make the same prediction. If lots of overlapping representations occur then that gist-like generalization will no longer rely on the hippocampus.

My research interests regarding memory specificity and generalization do not make an argument for either of these models. Both models suggest a potential trade-off between generalized memory schemas and more detail-oriented memories. This is because both theories focus on the involvement of the hippocampus in memory over a long period of time. My work investigating memory specificity and generalization did not address the hippocampus' role and as such does not support one theory over the other.

Source Theory

Source memory theory refers to how information is recalled (Banks, 2000; Glisky et al., 1995). A source memory is one that is recalled along with other associated context such as when or where the information was learned. After a memory has been consolidated, retrieval isn't always a perfect process. Sometimes a memory can be recalled, accompanied with a vague sense of familiarity. In the case of my dissertation projects, source memory would provide an example of memory specificity. When distinguishing between highly similar old and new stimuli, recalling the source of the original memory can help you make the correct selection. My proposed experiments rely heavily upon participants' ability to recognize specific stimuli studied during the learning portion of the experiments. Thus, recalling the source information surrounding that particular stimuli will improve a participant's accuracy although they are never asked to recall contextual information. Thus, my experiments didn't directly address source memory.

Prototype Theory

Memory generalization as defined thus far refers to multiple instances that overlap in some way as to create a gist representation. One could also describe this generalization as a form of category learning. There are two primary competing theories of categorization- prototype theory and exemplar theory. In the framework of prototype theory, categories are represented by prototypical gist constructions that are created from multiple, overlapping experiences with items of a category (Homa et al., 1979; Posner & Keele, 1968; Reed, 1972; Rosch & Mervis, 1975; Smith & Minda, 2002). New items are categorized by comparing them to the prototypes of existing categories and determining which prototype that item most matches. This theory rests on maintaining an abstracted schema or gist, rather than requiring memory of specific instances within a category.

Exemplar Theory

The second major theory to explain how categories are formed and how new items are compared against these categories is exemplar theory (Medin & Schaffer, 1978). Exemplar theory describes a process of comparing an item or object's features to a set of categories such that the item's category is decided by which category set of features it best matches (Medin & Smith, 1984). Exemplar theory, unlike prototype theory, is better at including items into categorical groups that can have very dissimilar items. For example, a penguin isn't a prototypical example of a bird and thus might be misclassified by prototype theory. Exemplar theory's use of shared multiple features within a category is much better equipped to deal with such irregularities (Nosofsky, 1991). However, at least one researcher studying category learning has argued that exemplar theory has been over-extended in its use in explaining human behavior, particularly when it comes to more abstract concepts (Murphy, 2016).

Despite decades of work, distinguishing between exemplar theory and prototype theory in human cognition remains a difficult task. This is partially due to the task demands of different experimental designs. For example, categorization of a single category (A, not A) and sorting of two or more categories (A, B) were often described in similar terms but seem to rely on different processes (Casale & Ashby, 2008; Little & Thulborn, 2006; D. Zeithamova et al., 2008). The perceptual and memory processes needed for categorization extend far beyond prototype and exemplar theory (for a review, see Ashby & Maddox, 2011).

While it is important to recognize the contributions of both the prototype and the exemplar theory in understanding category learning and perhaps memory gist, the experiments described here don't fully discriminate between the two. Only in chapter three do I provide weak evidence for the prototype theory over exemplar theory or other competing theories. This stands in opposition to the conceptualization that participants could improve on memory generalization because they have memorized enough specific exemplars to which they are comparing new items. Chapter three serves as the main evidence disambiguating these two possibilities and further discussion of these theories can be found there.

Memory Consolidation and Sleep

Memory Specificity and Abstraction

While researchers often focus on how to improve memory, 'good memory' in itself is a value judgment based on the situation. In some cases, it's useful to remember all the details surrounding an incident. If you're cycling home one day and are almost hit by a car, it's useful to remember every detail about the event in order to file a police report. But sometimes it's actually more useful to combine multiple individual instances to create a generalized schema, such as recalling what factors to look out for when biking with cars (like someone parallel parking or turning right). You don't need to remember every instance of a person turning right while you were passing, but you should recognize if a car is starting to turn so that you slow down for them to pass (and avoid the collision).

Memory specificity requires that a person recalls a memory with precision and is able to correctly reject items that weren't previously studied. One proposed neural mechanism for distinguishing highly similar stimuli is called pattern separation (Yassa & Stark, 2011).

According to this mechanism, correctly discriminating studied old items from similar items (sometimes called lures) is possible because of sparse encoding in the hippocampus which allows for similar items to be encoded in a non-overlapping manner (Bakker et al., 2008; O'Reilly & Norman, 2002). For any given stimulus, the hippocampus sparsely encodes the memory such that there is little to no overlap between the neurons used for one memory representation with other representations. This form of encoding allows for stimuli that are extremely similar to still be differentiated, assuming the sparse encodings (i.e., patterns) are separated and don't overlap- hence the name pattern separation. If engagement of the hippocampus through sparse encoding leads to increased accuracy on distinguishing highly similar memories, then memory specificity might be maintained through these neural patterns.

While much research has focused on how sleep boosts memory in terms of overall strength or maintaining specificity, less is currently known about how memory for generalizations or abstract knowledge is processed during sleep (Hanert et al., 2017; Marshall & Born, 2007; Saletin et al., 2016). One study examining sleep deprivation found that differences in the structures of hippocampal subfields was associated with a participants' quality of non-REM sleep (Saletin et al., 2016). Furthermore, the quality of non-REM sleep was linked with later successful recall of a previous memory task and thus demonstrates the importance, and individual variability, of sleep on memory. Another popular method of studying sleep's contribution to memory is to compare between two groups, sleep and wake, after a delay. A more recent study using this design tested pattern separation and found that participants who slept were better able to discriminate among highly similar items while those who remained awake declined in performance (Hanert et al., 2017).

Sleep provides one potential period during which communication between the hippocampus and neocortex can occur, thus allowing for generalization (Inostroza & Born, 2013). Through repetitive neural reactivations, sleep might allow for abstraction of underlying rules particularly across longer periods of time (Graveline & Wamsley, 2017a; Lutz et al., 2017). Researchers used the dot pattern classification task to test participants' category knowledge and recognition of old or new patterns after a 12-hour delay that included either sleep or wake (Graveline & Wamsley, 2017a). They found that the sleep group maintained better performance for old exemplars, new exemplars, and recognized prototypes that these patterns were abstracted from better than the wake group. Another 12-hour delay experiment used a visual version of the DRM paradigm to test participants' gist memory and their ability to recognize studied shapes (Lutz et al., 2017). While this study didn't show any differences between the sleep and wake groups after a 12-hour delay, a follow-up a year later showed a difference between the two groups such that the sleep group showed better gist memory than the wake group. Both of these experiments demonstrate that while generalization can be an outcome of a sleep period, the time scale of this generalization isn't well understood.

Spreading activation across the neocortex could be another explanation for how generalized memories are created. Spreading neural activity might allow for physically distant networks to interact and thus allow for reorganization of memory traces such as demonstrated in a 2009 experiment using the remote associates task (Cai et al., 2009). The researchers found increased performance on the task only after sleep compared to an equally long period of incubation and argue that this demonstrates that interactions between different neural networks during sleep, REM in particular, allows for this reorganization. During REM sleep, levels of acetylcholine are high in the hippocampus, which suppresses feedback from the hippocampus to the neocortex, while decreased levels of acetylcholine in the neocortex could promote spreading activation without hippocampal input (Hasselmo, 1999). Researchers argue that this process is why REM sleep in particular has been shown to increase creative problem solving; spreading activation helps the brain create an abstract memory representation though new connections between distant cortical regions (Boyce et al., 2017; Cai et al., 2009; Kim et al., 2017).

Endogenous memory replay during sleep

Memory replay was first recognized in rat models when researchers observed neural replay of place cells activated during a previous maze task (Pavlides & Winson, 1989). A later study found the same pattern of hippocampal place cells that were active in exploring rats during wake also fired in the same pattern during subsequent sleep at an increased rate (Wilson & McNaughton, 1994). Place cells that were inactive during the exploring period prior to sleep didn't show this increased replay, indicating that the replay of the related place cells is tied to memory consolidation of that exploration. More recent research has demonstrated this replay during sleep is specific to place cells that were previously activated during a maze exploration task and that this activity includes communication between the hippocampus and neocortex (Ji & Wilson, 2007). This initial animal research laid the groundwork for exploring memory replay and reactivation in humans.

The active system consolidation hypothesis is one of the more influential theories to explain memory replay during sleep (Diekelmann & Born, 2010). Some proponents of the active system consolidation model argue that memories are consolidated in a two-part process, dependent on two different stages of sleep. During the initial slow-wave sleep (SWS) that dominates the first part of the night, neural patterns are reactivated and representations that are important are tagged for later consolidation during the later stage of REM sleep.

However, sleep is composed of more stages and physiological events than the general dichotomy of SWS and REM sleep. Within the slow oscillations (less than 1 Hz) that give slow-wave sleep its name, there are nested oscillations of thalamocortical sleep spindles (11-16 Hz) and sharp-wave ripples that facilitate communication between the hippocampus, thalamus, and neocortex. The communication among these regions is thought to explain how memories are redistributed from the hippocampus to more permanent neocortical storage sites.

In more recent years, researchers posit that specific neural oscillations facilitate memory consolidation, although this literature isn't clear as to which oscillations are most important (Antony, Cheng, Brooks, Paller, & Norman, 2018; Cairney, Ashton, Roshchupkina, & Sobczak, 2015; Mölle, Bergmann, Marshall, & Born, 2011). For example in one sleep cueing study, beta power (16-30 Hz) after a cue predicted later forgetting while sigma power (11-16 Hz) predicted better retention after sleep (Antony, Cheng, et al., 2018). Other work has found that spindles after a sleep cue is associated with improved memory performance after sleep (Cairney et al., 2015). And a third study recording EEG with a high-density array in Germany found that specifically fast sleep spindles that occur during slow oscillation up-states are correlated with later memory performance (Mölle et al., 2011).

Slow-wave sleep

In order to better understand sleep consolidation, we can examine the frequencies, amplitudes, and duration of neural oscillations through electroencephalography (EEG) and how those states relate to later memory performance. Early research on memory and sleep focused on the contrast between SWS and REM sleep (for a review, see Stickgold & Walker, 2005). Slow-wave sleep measures such as delta power, time spent in SWS, and percentage of total sleep time spent in SWS have been linked to increases in declarative memory and spatial memory in particular (Gais & Born, 2004; Hanert et al., 2017; Marshall & Born, 2007; Tamminen et al., 2010).

I hypothesize that SWS is important not just for memory consolidation, but also for forming generalizations across multiple memories. Slow wave oscillations are created by many neurons across the cortex firing and going quiet at approximately the same time. This could lead to activation of multiple, overlapping memory traces at the same time and allow for generalized rules or schemas to form. In a recent experiment, researchers found that participants who had SWS after a challenging videogame were more likely to solve the puzzles in the game than those who did not get SWS (Beijamini et al., 2014). This could be due to SWS creating overlapping representations of the participants' experience in a way that allowed them to realize the correct solution that they missed prior to the nap. Other studies have also found improved problem solving after a night of sleep, although these studies often don't include EEG monitoring and so we can't be certain of SWS's contribution (Ellenbogen, Payne, et al., 2006; Lau et al., 2011; Sio et al., 2013).

And in some cases, SWS might allow memories to generalize so much that important details are lost. If the overlapping memories create a generalized schema without unique details, it could lead a participant to assume other categorically-related items were studied and thus cause false recollections. In a 2009 experiment, Payne and fellow researchers tested how false memories change across wake or sleep. They tested both a short nap delay as well as an overnight design. In both studies of this experiment, the researchers found that more time spent in SWS correlated with decreased recall which could be interpreted as the participants' memory becoming more generalized and thus less able to recall specific items from the list (Payne et al., 2009).

If there is a trade-off between memory specificity and generalization, that should entail a unidirectional relationship that details about a memory to fade while a gist representation remains. Researchers using Chinese characters as stimuli found that trade-off in a 2011 study (Lau et al., 2011). Comparing across sleep and wake, Lau and colleagues found that the sleep group performed worse at identifying old characters they'd studied. However, the sleep group was better than the wake group at identifying the meaning of radicals (a part of the Chinese character that refers to a general concept) and guessing the meaning of new characters that included radicals that they'd studied. This seems to indicate that a trade-off could exist, although this study didn't record sleep physiology and cannot make arguments directly about SWS. And if it were true that SWS leads to a more generalized memory representation, how are we able to retain memory specificity after a night's sleep. One possible candidate is sleep spindles.

Sleep spindles

Sleep spindles are rapid bursts of neural activity at 11-16 Hz, and are correlated with improved performance on memory tasks (Cairney et al., 2015; Mednick et al., 2013; Molle, 2006). Sleep spindles are traditionally split into fast spindles that are located centro-parietally and slow spindles that are fronto-central (Mölle et al., 2011). Because spindles tend to occur during slow-wave sleep, previous research has focused on how sleep spindles relate to declarative memory (Cox et al., 2012; Schabus et al., 2004). Rodent research has shown that sleep spindles often serve as envelopes for high frequency sharp-wave ripples (SWR) from the hippocampus (Molle, 2006). The SWRs imbedded in sleep spindles (that tend to ride on slow wave activity) provide a framework for communication between distant brain regions and therefore may promote memory consolidation.

In one 2013 experiment, researchers used click sounds in an attempt to enhance slow oscillations that also increased the number of both fast and slow sleep spindles. The researchers found that the number of induced sleep spindles was correlated with memory performance on a word pair task (Ngo et al., 2013). In a more recent iteration of this experiment, the researchers used transcranial alternating current stimulation (tACS) to boost sleep spindles. They found that memory improved in the stimulation condition and was correlated with sleep spindles, but not slow oscillations (Lustenberger et al., 2016).

As previously mentioned, sleep spindles can be differentiated into fast and slow spindles with many recent experiments pointing particularly to the importance of fast spindles. Some researchers posit that fast and slow spindles have different neural generators, occur at different points during slow oscillations, and that fast spindles are linked to memory consolidation (Mölle et al., 2011). In a sleep cueing experiment, the researchers found that the number of fast spindles predicted later explicit recall and improved reaction time for a motor task (Cousins et al., 2014). Fast sleep spindles during the up-state of slow oscillations also correlated with better performance on remembering word-scene pairs (Cairney et al., 2018). This evidence points to the possibility that fast spindles, which potentially have hippocampal sharp-wave ripples nested within, help protect memory consolidation and memory specificity in particular.

REM Sleep

Compared to SWS and associated sleep spindles, REM sleep has typically benefitted motor and emotional memory (Nishida et al., 2009; van der Helm et al., 2011). In the 2009 study, participants studied a collection of emotional and neutral pictures and then either had an afternoon nap or remained awake (Nishida et al., 2009). As might be expected, the sleep group showed better overall memory for both categories of images. When the researchers looked at the sleep physiology, they found that improved memory for the emotionally salient images was correlated with the amount of time spent in REM sleep across the nap. Other researchers investigating amygdala activity with regards to emotional stimuli also found that time spent in REM sleep was associated with later decreased amygdala activity and reduced subjective emotionality the next day (van der Helm et al., 2011). Motor learning tasks such as the serial reaction time task have also shown greater brain activity using PET during REM sleep in those who were trained on the task than participants who weren't trained (Maquet et al., 2000).

This classic view of REM sleep has shifted in recent years, with some creativity researchers positing that REM might not benefit a type of memory but rather a type of memory process. Researchers studying insight and creative problem solving have found REM to be correlated with making more insights and participants becoming explicitly aware of implicit rules (Alger & Payne, 2016; Kirov, Kolev, Verleger, & Yordanova, 2015; Stickgold, Scott, Rittenhouse, & Hobson, 1999; Tamminen, Lambon Ralph, & Lewis, 2017). One example of REM-related insight was a 2016 study where participants encoded AB and BC pairings of images (Alger & Payne, 2016). Participants were tested on their inferential knowledge (AC) after a nap and the researchers found that time spent in REM sleep correlated with better performance on these inferences. REM sleep also appears to be important in rule learning as an experiment using the serial reaction time task revealed that the more REM-NREM transitions during sleep correlated with more explicit knowledge of the hidden rule (Kirov et al., 2015).

So instead of stating that REM helps consolidate a type of memory, REM might actually be processing memories in a different way that allows for novel connections with existing knowledge and general reorganization of information (Boyce et al., 2017; Hasselmo, 1999). Additionally, REM sleep might also provide a period of secure memory storage, as has been suggested in at least one study (Batterink, Westerberg, & Paller, 2017). In this experiment, participants learned a set of new vocabulary words, half of which were cued during SWS. While there was no direct relationship between cueing and later performance, participants with greater amounts of REM showed a significant benefit for cued relative to uncued words. However, those who had little or no REM showed the opposite effect such that uncued words were better remembered than the cued ones. The researchers explain that their findings indicate that REM after SWS may be necessary for consolidating highly integrated memories.

REM sleep is unique among sleep stages not only because of definitional rapid eye movements, but also because it is a period of heightened autonomic nervous system (ANS) arousal (Whitehurst et al., 2016). Mednick and several other researchers argue that this ANS activity is linked to sleep consolidation. In an experiment using the remote associates task (RAT), Whitehurst and colleagues (2016) found that heart rate variability, which is a biological marker of ANS arousal, during REM sleep predicted improved post-sleep performance. This study, and a subsequent study using the RAT and a word analogy task both promote the importance of REM sleep for memory consolidation that leads to improved creativity and thus performance on these tasks (Cai et al., 2009).

While REM might be important for memory consolidation, at least one experiment has found that just a few minutes of REM doesn’t benefit later memory performance (Schapiro, McDevitt, et al., 2017). In these two experiments, participants learned three categories of satellite images. In the first experiment, researchers found that sleep compared to wake preserved both the specifics of previously-studied satellites and generalization of identifying the correct category for never-studied satellites. However, because participants either learned early in the morning or late in the evening, there were large circadian effects that muddied the results. The researchers then ran a second experiment where participants began the task at the same time of day and either napped or stayed awake in order to avoid these circadian influences. In this study, the nap group's memory improvement for low-frequency shared features (infrequent generalizable features) on the satellites was correlated with REM sleep. However, only participants who experienced a large amount of REM sleep showed this benefit. Thus, the story on REM's contribution to memory specificity and generalization remains murky at best. While we will investigate REM's role in memory consolidation in these proposed experiments, it is not the main focus of our investigation. This is because longer periods of uninterrupted REM tend to occur in the second half of nighttime sleep, but most afternoon nap studies consist of mostly SWS sleep stage. Thus, while we have some data to speak to this question, the proposed experiments weren't meant to test REM sleep specifically.

Targeted Memory Reactivation

Although memory replay can occur during wake or sleep, the unique neural circumstances of sleep allow replay to occur without competing with other on-going mental processes. Memory replay occurs endogenously during sleep, but can be directed using targeted memory reactivation (TMR): an experimental tool that allows researchers to harness replay for specific items during sleep. During a learning phase, participants are exposed to an external stimulus such as an odor or sound. Across the course of learning, the participant associates the stimulus with an individual item or set of items so that when the person is re-exposed to the stimulus during sleep, it will reactivate that memory trace (Oudiette & Paller, 2013; Schouten, Pereira, Tops, & Louzada, 2017). Although sound cues can be administered during at any point during sleep, early work demonstrated that SWS is a more effective time for spatial memory reactivation than REM sleep (Rasch, Buchel, Gais, & Born, 2007). Further research demonstrated that timing of the sound cues in regards to ongoing oscillations or sleep spindles also affects the cue's efficacy (Antony, Piloto, et al., 2018; Batterink et al., 2016). Prior TMR studies have often not taken into account ongoing slow oscillations when the cues are presented. In one experiment, the researchers found that cues played during the up-states of the slow oscillations showed a larger effect than those played during the down-states (Batterink et al., 2016). In more recent work, the timing of sleep spindles after a cue have also found to impact later memory performance. When looking at a period of three to six seconds after a sleep spindle has occurred, researchers found that cues presented just after this refractory period were associated with better memory consolidation than cues played during the spindle itself (Antony, Piloto, et al., 2018).

Using sounds as memory cues presents an added benefit over odors in that they are easier to administer and are easier to end if a sleep participant starts to awaken. Each unique sound can then be tested individually for improvement between encoding and the post-sleep test (Rudoy et al., 2009). Because multiple sound cues can be used within a single experiment, researchers can choose to cue half of the items and withhold the other half such that each participants' uncued items act as a control. This allows for a stronger experimental design with within-subjects comparisons.

The type of memory replay elicited by TMR has previously centered on hippocampal-based memory required for spatial tasks. However, other experimental designs have been employed to examine a broader array of memory types including motor, sensory, and language learning (Cousins et al., 2014; Honma et al., 2016; Schreiner, Göldi, et al., 2015).

Memory precision and integration in TMR

One of the more interesting uses of TMR is to employ sleep-cues to improve foreign language learning. In fact, several experiments have tested TMR for exactly this use. A prior experiment based on vocabulary learning of Dutch words by German speakers found TMR improved recall but not recognition for individual words (Schreiner, Göldi, et al., 2015; Schreiner & Rasch, 2015). The researchers played the Dutch words during non-REM (stage 2 and SWS) and measured later performance. They found that cueing improved later memory for the German translation of that word; and that an increased number of frontal slow waves correlated with these results. These findings could indicate that slow oscillations are important for consolidating within semantic language networks in the brain.

Unlike experiments in learning foreign language words, real-world language learning requires both vocabulary and grammar. One TMR study dealing with grammatical structure for language learning did find a positive effect for cueing, but only if the nap period also included REM sleep (Batterink et al., 2017). In a language labeling experiment, participants were required to learn made-up names for either familiar or totally unfamiliar objects. The researchers found that only the familiar objects with made-up names benefitted from TMR, while the unfamiliar objects showed no cueing benefit (Groch et al., 2017).

In terms of memory integration, some researchers have used olfactory TMR cues played during SWS to examine brain activity after a fear-learning condition and a nap. The researchers found less engagement of the hippocampus during test, perhaps indicating that cueing causes memories to be less dependent on the hippocampus (Hauner et al., 2013).

TMR memory consolidation and REM sleep

Some researchers have hypothesized that TMR cues during SWS are only half the story. In a 2017 experiment, participants learned a set of new words and their meanings which were then cued during SWS. After the nap, researchers found that TMR didn't affect later memory performance, but that the number of minutes spent in REM did correlate with competition between cued and uncued words (Tamminen et al., 2017). The researchers found that while sleep protected against forgetting words in general, REM sleep correlated with increased lexical competition between cued words and similar sounding non-words. The researchers explained these findings as evidence that SWS might be 'tagging memories' for consolidation which are later integrated during REM sleep. Indeed, researchers studying rats have found that there is neural replay of place cells during cued SWS, but there is also replay occurring during REM (Bendor & Wilson, 2012).

Thus, researchers are not arguing that REM sleep alone is enough to consolidate memory. More accurately, some researchers hypothesize that there is a two-part system where cueing during SWS allows for general consolidation while REM sleep is important for integration and the formation of unexpected connections between existing memory (Lewis et al., 2018). Such a model would predict that SWS would improve memory specificity, that cueing would enhance this process, and that REM sleep would improve memory generalization. While this is a plausible theory, the field still hasn't addressed the open question of how cueing might differently affect gist-based memory as compared to memory specificity. For a full review of TMR's role in memory specificity and generalization, see chapter two.

Completed Studies

In order to test sleep's contribution to memory specificity and generalization, I used two core experiments in an attempt to provide converging evidence. In the first experiment, I contrasted equally long periods of sleep and wake between participants that were trained on a series of outdoor scenes. Because of circadian effects, I followed-up this design with a second design where all participants came to the lab at the same time of day. They were then either randomly selected for an afternoon nap or were asked to remain awake and listen to a podcast while we recorded brain activity. I hypothesized that the sleep group would generally outperform the wake group after the delay in both studies. Furthermore, I hypothesized that specific sleep physiology measured during the second study would correlate with later performance such that fast spindles would be associated with better memory specificity while slow oscillations would be associated with better generalization.

In my second study, I used TMR on participants while they took an afternoon nap after learning a series of paintings created by six different artists. My aim was to test whether TMR altered sleep consolidation tradeoffs between memory specificity and generalization. I hypothesized that TMR would improve both memory for generalization as well as memory for specific paintings. Additionally, I hypothesized that the cueing effect for generalization would be larger than that for remembering specific paintings.

CHAPTER 2

Examining sleep’s role in memory generalization and specificity through the lends of targeted memory reactivation

Two vital memory functions — remembering specific experiences and generalizing across many experiences — are in tension with each other. In the complementary-learning-systems model, the hippocampus allows for fast learning of unique episodic memories while the cortex slowly extracts regularities from overlapping representations. Whereas episodic memories undergo consolidation over protracted time periods, many questions remain about how memory generalization evolves over time. Sleep’s role in consolidating individual memories has been convincingly demonstrated using targeted memory reactivation, a method whereby memories can be selectively strengthened through the unobtrusive presentation of learning-related stimuli during sleep. In this review, we argue that targeted memory reactivation can help advance understanding of memory transformation and the contrast between specificity and generalization.

Introduction

Declarative memory encompasses the ability to remember specific experiences and to generalize across multiple experiences (i.e. event and fact learning). For example, if you went to an aquarium you would spot an unusually shaped ocean sunfish and commit its specific, unique features to memory. Yet, perhaps you have seen many kinds of sea creatures while scuba diving, so the sunfish could strike you as similar, but not quite identical, to a filefish. A complementary-learning-systems orientation (McClelland & O’Reilly, 1995; O’Reilly & Norman, 2002) would posit that the hippocampus quickly encodes individual memories (a specific sighting of a unique sea creature) whereas the cortex slowly extracts regularities from overlapping representations (generalizing based on cumulative experiences with many sea creatures). Does sleep similarly impact these two memory functions, or are specificity and memory generalization afforded special and distinct processing during sleep?

Memory consolidation refers to the process of stabilizing and integrating memories within the brain. Consolidation may largely depend on the reactivation of previously acquired memories (Paller et al., 2020). Although memory reactivation can occur during wake or sleep, the unique circumstances of sleep may be advantageous in allowing reactivation to occur with minimal competition from other information processing. Additionally, some researchers argue that sleep provides an opportunity for communication between the hippocampus and the neocortex to allow for decontextualization in conjunction with a transfer of information to extrahippocampal regions (Inostroza & Born, 2013). Generalization could be considered a consequence of transformation whereby memories become less hippocampal-dependent and rely more on distributed traces across the neocortex. Others argue that the reactivation of overlapping memories during sleep strengthens the shared features and leads to generalization (Lewis & Durrant, 2011).

Considering the current models of sleep's role in memory, it is unclear if detailed memory representations are in direct conflict with generalized gist representations or whether the two coexist (Figure 2.1). Strengthening of overlapping memory traces may result in a loss of unique detailed features. Alternatively, generalization and memory for specifics could coexist such that shared features across representations are strengthened to create a new network without losing unique features of episodic memories.

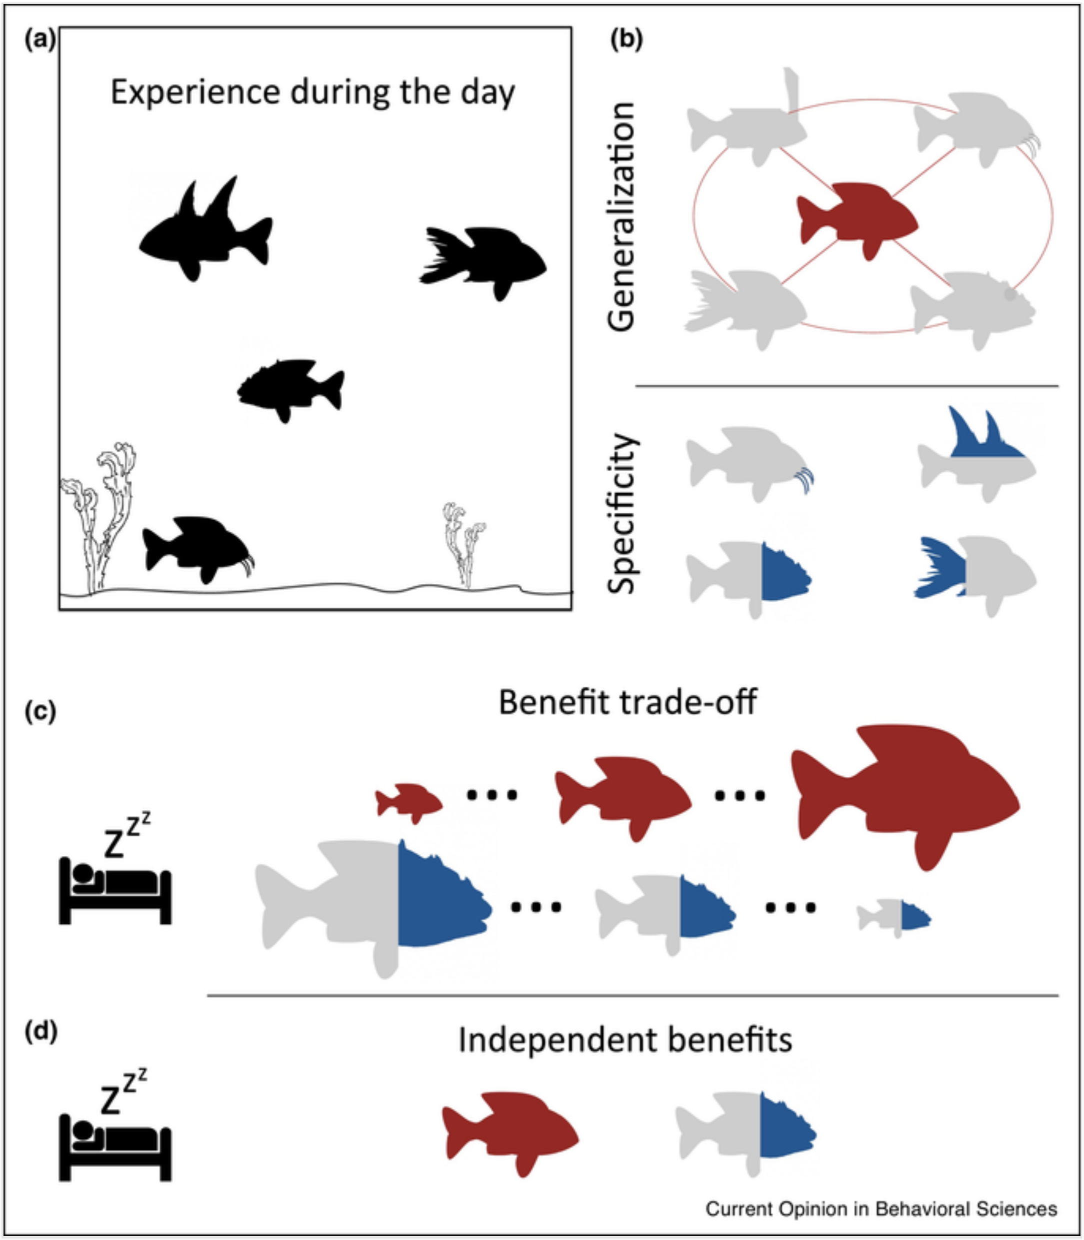

Figure 2.1 Memory consolidation for specific details and generalization over sleep. This diagram provides a schematic illustration of how memory consolidation may occur over sleep. **(a)** In this example, a person sees a set of different fish. **(b)** Memory of the experience evolves over time in two ways. In generalization (top), shared features among the fish contribute to a gist representation or abstracted prototype (represented in red). Memory for unique features of each fish (represented in blue) may also be consolidated to support specific remembering (bottom). Generalization and specificity may both be influenced by memory consolidation during sleep. **(c)** One hypothesis is that memory consolidation could help generalization over sleep (schematized as increasing size of red prototype), at the cost of memory for specifics (schematized as decreasing size of blue specifics), or *vice-versa* **(d)** Alternatively, sleep

consolidation could preserve memory for specifics and generalized memories independently of each other.

Sleep's impact on declarative memory for specifics

Sleep's protective role against forgetting has been recognized as far back as the 1920s (Jenkins & Dallenbach, 1924), but this idea did not catch on quickly in memory research. Recent emphasis on the period of sleep known as slow-wave sleep (SWS) coincided with a surge of interest in this topic (Gais & Born, 2004). A landmark observation was that rodent hippocampal place cells display firing patterns during sleep that recapitulate those in previous wake exploration (Foster, 2017; Pavlides & Winson, 1989). In particular, replay in the rodent hippocampus was shown to be temporally coordinated with memory reactivation in the visual cortex during SWS (Ji & Wilson, 2007).

Research into human memory also supports the idea that sleep, and SWS in particular, is important for memory for specifics. This form of memory can be considered in relation to the neural process of pattern separation, whereby memory representations can be distinctive when they overlap minimally with each other (Yassa & Stark, 2011). Distinct representations allow for the successful retrieval of specific details and discriminating between similar representations in the face of potential interference. To study pattern separation, many researchers employ the mnemonic-similarity task, in which participants are first exposed to a series of objects and categorize them as either indoor or outdoor objects, followed by a recognition test with novel objects, old objects, and highly similar objects (Bakker et al., 2008; Doxey et al., 2018; Hanert et al., 2017; Stark et al., 2013). Accurate recognition of old objects with low false-alarm rates, particularly for the highly similar objects, is indicative of high memory specificity, due presumably to effective pattern separation. Using such a task, researchers compared a 12-hour delay including either sleep or wake (Doxey et al., 2018). They found that sleep preserved memory specificity more than when participants remained awake.

These findings were corroborated by results from another study with a 9-hour retention interval during the day or overnight with electroencephalographic (EEG) monitoring (Hanert et al., 2017). Hanert *et al.* found that high specificity in recognition was positively correlated with two sleep physiology signals, slow oscillations (0.5 - 1 Hz) and sleep spindles. Sleep spindles are rapid bursts of neural activity at 11-16 Hz, which can be observed during SWS (Cox et al., 2012; Schabus et al., 2004) and are correlated with improved performance on declarative memory tasks (Antony, Schönauer, et al., 2018; Cairney et al., 2015; Mednick et al., 2013). These electrophysiological findings were taken as evidence for sleeps' ability to improve hippocampal representations and enhance performance on highly specific memory tasks.

Memory generalization during sleep

There are many ways to study the process of generalization (Lerner & Gluck, 2019). For example, investigators have used procedures in which participants gradually learned to extract relationships among various abstract stimuli, to solve a puzzle based on repeated mathematical procedures, to produce a word that fits multiple constraints, or various other language tasks.

In a 2007 experiment, Ellenbogen *et al.* presented participants with pairs of abstract images that fit into a complex hierarchical structure, followed by a memory test with new pairs (Ellenbogen et al., 2007). Relational knowledge was assessed by determining whether participants inferred A > C (where '>' means 'should be selected over') after learning of A > B and B > C, with no prior exposure to the A/C pair. This generalized knowledge for the underlying structure of stimuli (also known as transitive inference) was stronger after a period of sleep compared to a period of being awake. However, measures of sleep or sleep-stage were not available, so it was unclear which aspects of sleep may have improved generalization.

A more recent study tested how sleep could improve insight for a hidden rule (Verleger et al., 2013). The researchers used the number-reduction task, which can be solved either by a sequence of simple math operations or by using a faster, hidden rule. The study replicated an earlier finding (Wagner et al., 2004) that those who slept were more likely to discover the rule. Also, beta power (17 - 25 Hz) predicted which participants would discover the rule, but only during SWS.

Although this study implicates SWS in memory generalization, in other relevant studies SWS has not consistently been the most important sleep stage. Rapid eye movement (REM) sleep, which tends to be more prevalent in the second half of the night, has been associated with integrating unassociated information for creative problem solving (Robert Stickgold & Walker, 2004). Using the remote-associates task, where participants see three cue words and must identify a fourth word linked with all three (i.e. crab, pie, and pine are all linked by apple), researchers compared wake participants with sleep groups that either did or did not enter REM sleep during an afternoon nap (Cai et al., 2009). Only the REM sleep group improved on the task by integrating unassociated information to find solutions. Although the authors concluded that REM sleep is important for integration, there may have been confounding factors that led some subjects to have more REM. Also, other studies have not always found REM sleep correlated with problem solving (Beijamini et al., 2014). In fact, a recent experiment using magic tricks and classic insight problems found no effect of sleep on problem solving at all (Schönauer et al., 2018).

Contrasting specificity and generalization during sleep

The experiments described thus far examined either memory for specifics or generalization. It could be advantageous to analyze both types of memory in tandem. This tactic was utilized in a study where learning of Chinese characters was followed by a 90-min afternoon nap or wake period (Lau et al., 2011). Recalling the meaning of studied characters was relatively worse after a delay with sleep compared to wake, but generalization was improved, operationalized in this study as the ability to recognize common symbols shared across characters (i.e. the symbol for woman is included in the characters for maid, princess, and nurse).

In another study also supporting the idea that sleep is preferentially beneficial for generalization (Schapiro, McDevitt, et al., 2017), participants viewed artificial images composed of shared features (defining categories) and unique features (identifiers for individuals). Sleep was associated with relatively better memory only for the former.

Researchers have also tested 15-month-olds' generalization abilities through exposure to triplets of spoken words in an artificial language with hidden dependencies between the first and final words (Gómez et al., 2006). Infants either napped or stayed awake for a 4-hour delay before a subsequent test. Whereas memory for the previously heard strings was evident in the wake group, an ability to abstract the grammatical relation and apply it to new strings of nonsense words was significant only in the sleep group. Together, these studies indicate that even a short period of sleep may preferentially influence generalized memory over memory for specifics.

In another study comparing sleep and wake groups across a 12-hour delay, participants completed a dot-pattern-classification task, where they categorized constellation-like images (Graveline & Wamsley, 2017a). Researchers found that overnight sleep improved categorization of new stimuli but had no effect on recognition of old stimuli.

In a longer-delay study that tested generalization in a visual-categorization paradigm, researchers tested participants after a 10-hour period containing sleep or wake, and then again 1 year later (Lutz et al., 2017). Participants who slept after learning were relatively better at recognition specificity at the 10-hour delay, but there was no difference between the groups on generalization. However, after a year delay, only the sleep group showed gist knowledge in their categorization performance, even though specific items were not remembered in either group.

A reasonable inference based on this evidence is that generalized memory and memory for specific details can be in competition with each other, with sleep enhancing memory generalization but not specific memories. However, this hypothesis is challenged by numerous studies showing that sleep improves memory for specifics (Doxey et al., 2018; Jegou et al., 2019). There are at least three possible explanations for this apparent conflict. First, the experimental designs above may have emphasized general features to the detriment of specifics during learning. Second, perhaps these paradigms lacked sufficient sensitivity to specificity effects. Finally, maybe both memory types benefit from sleep, depending on different sleep stages or sleep physiology.

All the above studies compared memory following wake versus sleep, sometimes emphasizing different sleep stages. However, these sleep and wake conditions vary in many ways, including alertness, degree of interference, and circadian factors. Accordingly, results from sleep-versus-wake designs are subject to interpretive limitations such that sleep’s role in promoting memory generalization and specificity has remained equivocal. However, alternative experimental strategies that avoid these shortcomings could be employed.

Future prospects using targeted memory reactivation

Memory reactivation during sleep can be biased using targeted memory reactivation (TMR), an experimental tool that allows researchers to direct reactivation for specific items during sleep. The procedure typically starts with a learning session including one or more special stimuli that are linked with aspects of learning. The same stimuli can then be presented during sleep — carefully to avoid producing arousal — in an attempt to reactivate memories (Oudiette & Paller, 2013; Schouten et al., 2017). A recent meta-analysis of the TMR literature to date demonstrated that such cues were effective during SWS and stage-2 sleep (Hu et al., 2020).

Many early studies using TMR focused on spatial memory where individual items are associated with cues and a specific location on a grid (Rasch et al., 2007; Rudoy et al., 2009). These designs are particularly powerful for detecting small, specific changes between cued and noncued items, as measured by error between an item’s studied location and the location recalled by the participant at test. TMR has been applied to investigate memory consolidation for a broad array of memory types (Cousins et al., 2014; Honma et al., 2016; Schreiner, Göldi, et al., 2015). Though most studies using TMR have focused on its impact on the specific aspects of remembering, TMR offers advantages for investigating sleep’s role in both memory generalization and specificity. Relying on within-subject comparisons between cued and noncued information, for example, avoids confounds with differential alertness, time of day, demand characteristics, and potential interference from waking experience. Although some evidence suggests that TMR-induced consolidation may differ in some respects from spontaneous reactivation (Klinzing et al., 2018), it remains a useful tool for understanding memory transformation over sleep.

One recent TMR experiment explored sleep's contribution to generalization by using lexical competition between words and nonwords (Tamminen et al., 2017). The researchers hypothesized that cueing during sleep would lead to better lexical integration of artificial words. Results showed no direct effect of TMR. However, for cued words, REM sleep was correlated with better integration of new information within existing knowledge. That is, participants took longer to make judgments for stimuli confusable with cued relative to the noncued words when they spent more time in REM sleep, indicating that TMR facilitated the extent to which words were embedded in the lexicon. While touching on generalization, this study does not reveal what aspect of generalization leads to competition nor does it test for specificity.

Rather than focusing on competition between specific words, researchers in another study used TMR to test grammatical rule abstraction (Batterink & Paller, 2017). Participants learned grammatical rules through a language task and then participated in a second, unrelated task. During an afternoon nap, participants were cued for either the grammar task or the unrelated task, and those who were exposed to the language phrases during sleep showed a relative gain in grammar learning. Whereas these findings reinforce the notion that memory reactivation can impact generalization, the performance criteria at learning were such that it was not straightforward to assess memory for specific phrases at test.

Insofar as sleep's role in the consolidation process for specific episodic memories and gist is not well understood, manipulating reactivation via TMR provides a suitable tool for seeking answers to these questions. Future TMR experiments should be designed with a variety of learning materials and procedures to test both integrated knowledge as well as specifics of the items studied. Ideally, these two features can be examined in parallel such that researchers can identify if they rely on different processes or represent a trade-off whereby specificity declines as generalized knowledge is gained (or *vice-versa*).

For example, analogical problem solving could be used to put generalization and memory specificity in opposition (Monaghan et al., 2015). Participants could attempt to solve a series of word problems with distinct features, half of which would be cued during sleep. Upon waking, participants would attempt to solve an analogous problem with different details, testing both for problem solving (generalization of the underlying structure) and details of the specific problems.

Additionally, better understanding of how TMR cues bias consolidation during sleep and whether reactivation is specific to the item or applies to a general context will improve experimental design for all studies going forward (Schechtman et al., 2019). Harnessing TMR to causally affect memory consolidation may prove crucial in teasing apart hypotheses regarding sleep's role in generalization and the possible trade-off between generalized and specific memories. We might discover that different sleep-physiology signals relate to how brain networks are engaged to preferentially yield either generalization or maintenance of details. Further studies might find that sleep promotes generalization at the expense of memory for specifics. Alternately, as schematized in Figure 2.1, sleep may both promote the integration of specific memories into broader schemas while also protecting the specific of those memories. Experimental approaches that take into account the dual role of consolidation in facilitating memory for specifics as well as generalization are crucial to improving our understanding of sleep’s role in memory.

CHAPTER 3

Experiment 1: Comparing sleep versus wake groups on memory for individual stimuli and categories of confusable terrains

Studies across a broad set of experiments have highlighted sleep as important not just for memory consolidation of specific instances, but also generalization (Stickgold & Walker, 2013). In one influential paper, the authors argued that memory consolidation processes during sleep might gradually cause memories to lose specific details such that long-term representations become more gist-like (Inostroza & Born, 2013). Although such gist extraction might be important for rule learning or generalizing behavior across similar circumstances, it remains unclear whether such abstraction occurs primarily during sleep or simply across extended periods of time. Additionally, researchers have yet to conclusively tackle whether abstraction, either across time or sleep, comes at the cost of specific memory.

Generalization across many instances to form a categorical representation has been studied in various ways, but one of the earliest variants is the prototype distortion task (Posner & Keele, 1968). In this experimental design, the stimuli are composed of dot patterns from a prototypical image where the resulting categorical exemplars are creating by jittering the dot locations from the original prototype. Although the design is rather straightforward, results from prototype distortion studies have been used to argue for both the exemplar and prototype models of categorization (Homa et al., 1979; Minda & Smith, 2002; Nosofsky, 1991).

Supporters of the exemplar theory argue that categorization is achieved by comparing a stimulus to previously seen specific, stored examples and determining which examples most closely match the new stimulus (Busemeyer et al., 1984; Medin & Schaffer, 1978). On the other hand, those who endorse the prototype theory contend that categorization is achieved by maintaining a gist memory representation abstracted from previous experiences (Homa et al., 1979; Minda & Smith, 2002; Posner & Keele, 1968). These two theories both account for aspects of categorization and occur during sleep or wake, but make different predictions about memory for specific instances. A prototype model would maintain categorization even if there was forgetting of specific stimuli while the exemplar model would expect categorization to drop if there was forgetting of specific instances.

One recent study used the prototype distortion task to compare between participants that either slept or remained awake over a 12-hour delay. The researchers found that sleep benefitted both studied and new exemplars compared to the wake group (Graveline & Wamsley, 2017b). The sleep group did not show a significant improvement on categorizing exemplars or identifying whether a dot pattern was new or old after a 12-hour delay. However, the sleep group did maintain performance on these measures while the wake group showed a significant decline in all measures of performance (categorizing and recognizing old or new exemplars) after the 12-hour delay. This study would seem to indicate that sleep consolidation follows the exemplar theory trajectory where both specifics and categorization are protected from significant forgetting.

Similarly, a more recent study used a satellite-identification task to again compare between twelve hours of daytime wakefulness compared to overnight sleep. The researchers showed that overnight sleep protected both memory for unique features of a ‘satellite’ image and improved memory for shared features across similar satellites while both measures declined over a day of wakefulness (Schapiro, McDevitt, et al., 2017).

However, other research has found the predicted tradeoffs between recognizing old stimuli and generalizing to new knowledge as would be expected in a prototype model. In a study where participants learned the meaning of Chinese characters, sleep during a 90-minute afternoon nap decreased their ability to recall the meaning of characters they had previously studied compared to the wake group. But the sleep group was better than the wake group at guessing the meaning of new, similar characters after the 90-minute delay (Lau et al., 2011). However, this study was conducted over a much shorter time period than the two described previously.

In fact, the time course for forgetting of specific instances compared to memory necessary for categorization have not been adequately addressed. Regardless of sleep, only a few studies have examined how memory changes over time. One early report examined category learning both immediately after learning and at a one-week follow up (Homa et al., 1981). The researchers found that the exemplar theory better explained performance at the immediate test but that the prototype theory was better supported at the longer delay. This could indicate that different memory systems drive behavior, depending how much the memory tests are separated in time. At least one functional MRI study supports this idea, finding neural evidence for a two-phase learning process across a distributed network (Little & Thulborn, 2006).

As of yet, there is no conclusive answer as to whether sleep protects or improves both memories for specifics as well as generalization in the form of categorization. Nor is it known how these two memory representations might change across time. Prior studies have relied on rather simplistic stimuli, limiting their explanatory value because they may not represent experiences in the natural world. In the current set of experiments, we trained participants on three categories of short videos of outdoor scenes and then tested performance on memory specificity and categorical knowledge after a delay containing either sleep or wake. In study 1, participants were split into either sleep or wake groups and were tested at 9am and 9pm respectively. In study 2, all participants arrived to the lab in the afternoon and after the learning phase were split into either the nap or wake group. This allowed us to compare specificity and categorization across long and short delays containing either sleep or wakefulness.

METHOD

Participants

A total of 90 members of the Northwestern University community who had no known history of neurological or sleep disorders participated. Participants either earned class credit or were paid for their time. We recruited 49 participants (30 female, 19 male) for the first study, but excluded the results from 9 of these participants: 5 did not complete the full experiment and 4 were consistently below chance (~33%) for the last third of the learning phase. Of the 40 participants included in the final sample 20 participants were randomized into the overnight sleep condition (2 left-handed, mean age = 21.8 yrs) while the remaining were placed in the wake condition (1 left-handed, mean age = 20.1 yrs).

Forty-one participants (26 female, 2 non-binary, 13 male) were recruited for the second study, but 4 participants were excluded because they were consistently below chance (~33%) for the last third of the learning phase. Thus, our final sample consisted of 37 participants who were all asked to go to bed later than usual the night before the study, wake up one hour earlier than usual in the morning, and to avoid caffeine the day of the experiment. Eighteen participants were randomized into the sleep group (all right-handed, mean age = 20.9 yrs) and 19 participants (all right-handed, mean age = 21 yrs) remained awake. The Northwestern University Institutional Review Board approved this procedure and informed consent was obtained from all participants.

Stimuli

Alpha Bravo Charlie

Figure 3.1. Still image examples of the three video categories with accompanying labels.

Using Unity game engine software, we rendered 60 video clips of outdoor scenes. Each 4-second clip was shot from the first-person perspective and depicted forward movement in a backwards "S" trajectory through each scene. Video clips varied based on 4 environmental dimensions (i.e. weather, time of day, hilliness, and amount of vegetation) that could be controlled during the rendering process. Each environmental dimension could vary in increments of 0.1 on a scale from 0 to 1. For example, values closer to 1 for the "Hilliness" dimension would render flat terrains, while values closer to 0 would render steeper terrains. For the "Time of Day" dimension, values close to 0 depicted sky and illumination features of morning, while values closer to 1 depicted sky and illumination features of sunset and nighttime.

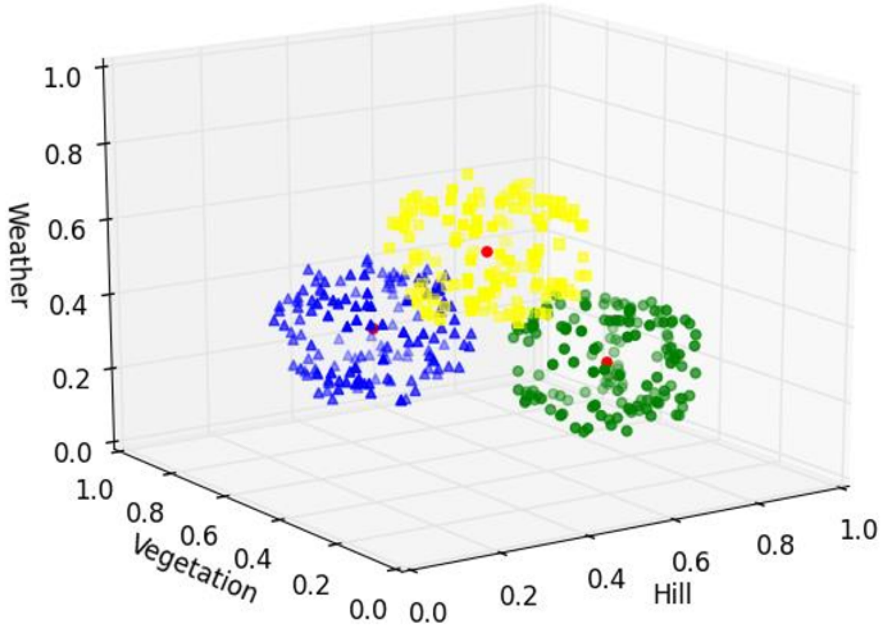

Figure 3.2. Representation of the three categories of terrain in a 3-dimensional space based on three of the four variables for producing the environment (time-of-day variable not shown here). The red dots represent the prototypes of each category while the points clustered around them represent the stimuli generated such that each video is equidistant from the prototype and none of the categories overlap. Alpha videos are represented as blue triangles, bravo videos are represented as green circles, and charlie videos are represented as yellow squares.

A “Prototypes-Exemplar” algorithm determined the chosen settings for each video clip by first plotting 3 “category prototype” point coordinates equidistant from one another in 4-dimensional space with each axis corresponding to an environment dimension scale, similar to selection of stimuli for the dot paradigm (Graveline & Wamsley, 2017b; Homa et al., 1979; Posner & Keele, 1968). Each “category prototype” point defined the exemplar category memberships to “Alpha”, “Bravo” and “Charlie” terrain types. Then, 20 exemplar point coordinates were generated to be a given distance from their respective category prototype point coordinate. The algorithmically chosen 4-D exemplar coordinates were then used to render each of the environmental dimensions of the video clip according to the specified values of each coordinate. Each generated video clip in a given category was equidistant in 4-D space from the initial prototype. This process ensured maximal balance between the similarity of exemplars in the same category membership and the necessity of maintaining uniqueness across all video clips. Participants were never exposed to video clips rendered from the prototype coordinates.

Procedure

Each participant gave written consent before beginning the experiment. In the first experiment, participants arrived for their first session at either 9am or 9pm, depending on whether they were randomly selected to participate in the sleep or wake group and were asked to complete the Morningness-Eveningness Questionnaire (MEQ) to determine their preferred circadian time. Additionally, during both visits participants completed the Stanford Sleepiness Scale (SSS) to record their drowsiness at the time of testing. These two measures were used to verify that group differences weren't due to drowsiness at test or individual differences according to circadian preference. Lastly, after the experiment was completed during the second session 12 hours after the initial session, participants filled out a general questionnaire to ask how they determined the categories for stimuli and to ask about general habits like caffeine consumption that have been known to affect circadian performance.

In the second experiment, participants arrived between noon and 2pm to begin the study. They also completed the MEQ before learning began, and answered the SSS both before and after the nap/wake delay. The procedure was identical to that of the 12-hour delay experiment except participants did not leave the lab between the two testing sessions. Instead, researchers took approximately 30 minutes to apply the Neuroscan electroencephalograph (EEG) cap before the participants were told in which group (nap versus wake) they had been randomly placed. The nap group slept in the testing chamber, with the futon armchair converted into a twin bed. The wake group remained seated in the testing chamber and listened to a 90-minute podcast from Gimlet Media's Reply All (episode #102, Long Distance).

EEG was recorded only during the 90-minute wake or nap phase from 21 scalp locations from the 10–20 system (Fpz, Fp1, Fp2, Cz, C3, C4, F3, F4, F7, F8, Pz, P3, P4, T3, T4, T5, T6, Oz, O1, O2) and both mastoids. Additional electrodes were placed on the face for recording vertical and horizontal electro-oculogram (EOG) and chin electromyogram (EMG). Electrodes were referenced to the left mastoid electrode and re-referenced offline to the average of the two mastoids. Impedances were brought down to 5 kΩ and voltage was sampled at 1000 Hz.

Learning. Participants were instructed that they would be viewing clips of outdoor scenes and that they were to determine to which category each clip belonged. Specifically, they were instructed,

> "You will be presented with short video feeds and asked to choose one of three possible categories identifying how the terrain differs, 'Alpha', 'Bravo', and 'Charlie'.
> Once you choose the terrain type, your computer will inform you if you were correct.
> If the selected terrain type was correct, a green 'Thumbs Up' will appear.
> If the selected terrain type was incorrect, a red 'Thumbs Down' with the name of the correct terrain type will appear.
> Note that if you aren't sure of which method to choose, please guess. In addition, the correct identification of terrain type depends on multiple features- ie, do not rely solely on the hilliness or cloud coverage alone.
> If you feel 'Alpha' is the correct method, press the key marked 'A'.
> If you feel 'Bravo' is the correct method, press the key marked 'B'.
> If you feel 'Charlie' is the correct method, press the key marked 'C'.

You will be able to take a break after each set of 30 video feeds. You will see a total of 30 videos each 9 times. All of this will be followed by a memory test over these specific videos so try to remember them as best you can. We know this is a very difficult task but do your best!"

After each 4-second clip, participants pressed "A," "B," or "C" to indicate whether the video was an Alpha, Bravo or Charlie. Participants had to respond within 5 seconds and were then provided feedback on whether they selected correctly and the correct category for the video if they selected the wrong category, lasting 2 seconds. Participants viewed 30 clips per round with 10 videos per category. The order of these clips was then randomized and displayed again. The 30 videos were repeatedly displayed for a total of 9 blocks (270 clip presentations) and were presented in a randomized order for each of the blocks. Participants did not know that there were only 30 unique videos shown during the learning portion. Participants took approximately 35 minutes to complete the learning portion.

Test 1 (T1). Directly after the learning session, participants were tested on their knowledge of both categorical information and whether they had seen a clip before (called "old/new"). Participants were instructed that they would see 60 total clips comprised of 30 videos repeated twice. Videos were tested twice in order to increase the number of stimuli without increasing the number of videos seen during learning. Half of the videos (15) were new, meaning they weren't studied during the learning portion, and half were old, indicating they were seen during learning. After each 4-second video, participants had to make an old/new judgment within 5 seconds, followed by a confidence rating (on a 1-3 scale) within 5 seconds. After which, participants were asked to identify if the clip was an Alpha, Bravo, or Charlie (5 second trial) and gave a confidence rating for that judgment (5 second trial). Participants had to respond to each question within 5 seconds and the entire pre-delay test section lasted approximately 10 minutes.

Test 2 (T2). The post-delay test was identical to the Test 1 but used a combination of the remaining 15 untested old videos as well as 15 totally new videos. Once again, the clips were each tested twice, meaning participants saw a total of 60 clips. This test lasted approximately 10 minutes occurred either 12 hours (study 1) or 2 hours (study 2) after Test 1 with either sleep or wakefulness over the intervening time.

Analyses

Behavioral Data. The primary measure of interest was how memory changed over time between the sleep and wake groups in both experiments. We used a two-way repeated measure ANOVA to compare the main effects of group and test session as well as the interaction between them. Group included two levels (sleep, wake) and test sessions consisted of two levels (T1, T2). The two-way repeated measure ANOVA was completed twice, once for the categorical test and once for the old/new specificity test. Post-hoc tests were used to determine whether results were due to time-of-day confounding factors rather than the group assignment. Specifically, we used a t-test to compare SSS scores between the sleep and wake groups to determine whether one group was drowsier during testing than that other. We also compared MEQ scores between the two groups using a t-test to ensure time-of-day preference didn’t vary between the groups.

Sleep Physiology in Study 2. Continuous EEG data were down-sampled to 128 Hz and filtered at 0.5–50 Hz using an infinite impulse response Butterworth filter. Sleep stages were formally identified using standard sleep scoring criteria. Wake EEG was also formally checked to verify that the wake participants remained alert for the duration of the 90-minute delay.

For the nap group, standard analyses of sleep oscillations were computed focusing on two clusters of interest (frontal using Fpz, Fp1, Fp2; central-parietal using Cz, Pz, C3, C4, P3, P4). A fast Fourier transform using a Hanning function and 5-s intervals was performed on NREM sleep epochs. We extracted mean power for delta (1–4 Hz) and sigma (12–15 Hz) bands. For slow-oscillation analyses, EEG was low-pass filtered at 3.5 Hz. Slow oscillations were detected by finding adjacent points in which the EEG signal voltage changed from positive to negative that were 0.5–2.0 s apart from each other, and when the maximum peak-to-peak amplitude between the two points was greater than 75 μV. Sleep spindles were automatically detected (Mölle et al., 2011) by first filtering EEG data between 11 and 16 Hz and calculating root mean squared (RMS) voltage using a sliding 200-ms window. A spindle was counted if the RMS crossed a threshold of 1.5 standard deviations of the signal and remained above the threshold for 0.5–3.0 s. We separately analyzed fast (> 13.5 Hz) and slow spindles (< 13.5 Hz) because they show different topographies, with fast spindles predominant at parietal and central locations and slow spindles at frontal locations.

Post Hoc Analyses Examining Group and Individual-Differences

Machine Learning. In order to verify if there were any differences between the sleep and wake groups, we ran a principal components analysis on participants' behavioral measures (such as confidence ratings, recognition accuracy, categorization, etc.) and selected 10 components from both the study 1 participants (n = 40) and the full set of participants of both studies (n = 77). These were used as input for a support vector machine (SVM) in python. We trained the SVM on all data except data from one subject that was held out, and tested whether we could correctly identify that one as belonging to either the sleep or wake group. We repeated this process, iterating though each participant before summarizing the results. We didn't find any components that reliably identified group membership above chance. We later attempted the same process using 20 components, but the SVM still wasn't able to identify group membership with any reliability and so those results aren't included here.

Explicit Categorization Rule Learning in Study 1. After completing the second test in Study 1, we asked participants to complete a general questionnaire including questions about caffeine consumption and their thoughts on the experiment. One of the questions asked was, 'How did you decide between the Alpha, Bravo, and Charlie categories?' We computed the number of rules that participants listed based on prior work from Brooke Feinstein in the Reber lab. Rules were defined as a combination of both a specific category name and a specific feature. For example, "Bravo looked like a sunset" would count as a rule while "I guessed" or "I used the camera angle" were not considered rules. Using numeric scores for each participant's rule count, we analyzed whether those with more rules performed differently on the task than those with fewer or no rules. We did not find any significant differences and so those results are not reported here.

STUDY 1 RESULTS AND DISCUSSION

Sleepiness survey and circadian preference. The sleepiness score from the SSS assessment did not differ between the sleep and wake conditions for T1 (mean sleep = 2.65; mean wake = 2.70; $t[38] = 0.14$, $p = 0.89$), but did differ at T2 (mean sleep = 2.65; mean wake = 1.85; $t[38] = -2.43$, $p = 0.02$), meaning the wake group was more alert at T2. This suggests that there were no alertness differences between the groups during the learning portion (T1) due to time of day, but differences did emerge after the delay. Additionally, the MEQ scores did not differ between the two groups (mean sleep = 41.95 i.e., no strong preference; mean wake = 42.95 i.e., no strong preference; *t*[38] = 0.46, *p* = 0.64), suggesting that there were no differences in circadian preference between the two.

Training. Participants acquired category knowledge over a series of 9 blocks, as shown in Figure 3.3. Collapsing across the last three blocks, average performance across all participants was 73.09% (SD = 16.49%). Due to a computer error, some data were missing in later blocks and have been removed from training analyses. We found there were no differences in learning between the two groups on any block, though most importantly they weren't significantly different at the final block as shown in figure 3.3 (mean sleep = 72.16%; mean wake = 75.21%; *t*[38] = -0.11, *p* = 0.91).

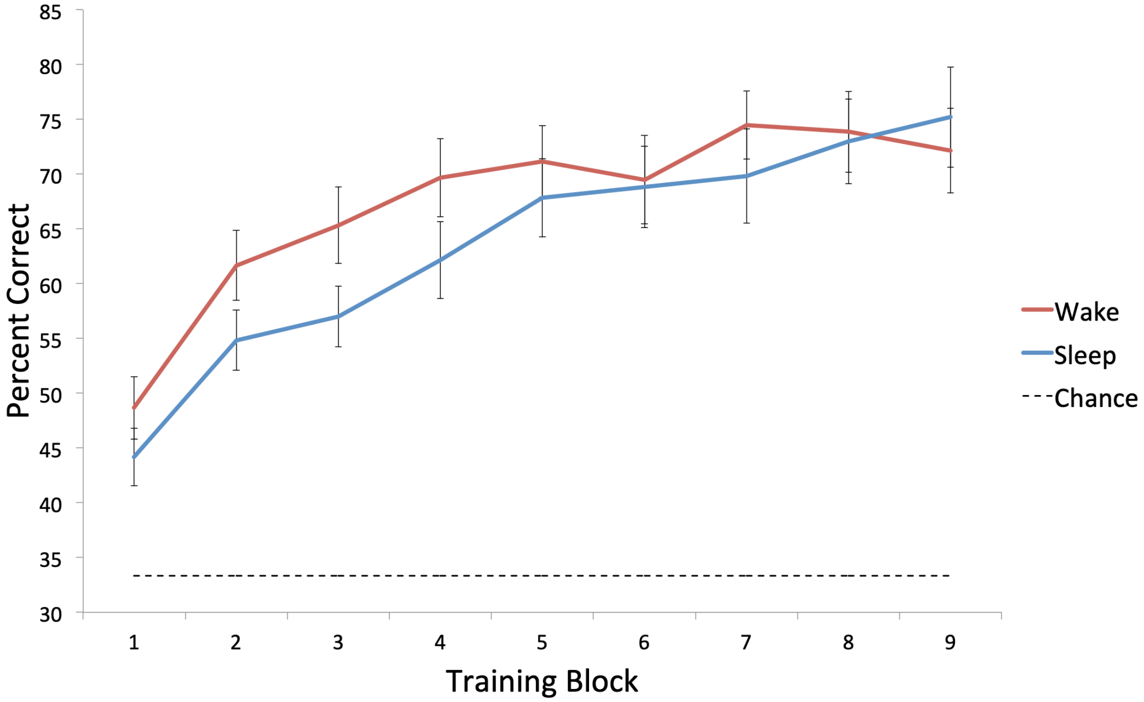

Figure 3.3. Line graph of average learning curves across the 9 training blocks for both the sleep and wake groups separately. The x-axis denotes the training block and the y-axis denotes the average percent correct for that block. The dashed line represents chance (33.33%).

Memory Specificity. On average, participants correctly selected whether the video was old or new with 57.31% accuracy (SD = 8.43%) and rated their confidence an average of 1.97 (on a 1 to 3 scale; SD = .32). Examining just items called new, participants on average were correct (hit rate) for 55.91% of the trials and false alarmed (called an old video new) on 29.83% of the trials, resulting in a d-prime of .75 at T1 (z-scored hit rate = .17, z-scored false alarm = -.60). After the delay, the average hit rate was 47.50% (z-scored value = -.06) and average false alarm of 41.17% (z-scored value = -.24), resulting in a d-prime of .17at T2. We found no difference in performance between the sleep and wake groups at T1 for recognition of whether a video was old or new (mean sleep = 62.75%; mean wake = 62.67%; $t[38] = -0.4$, $p = .97$). Upon further breaking this down, the sleep and wake groups weren't significantly different at recognizing only new videos ($p = .90$) or only old ($p = .89$). Finally, we analyzed participants' confidence at T1 and found no difference between the groups ($p = .63$).

To assess the change in recognition accuracy from T1 to T2, we used a two-way repeated measure ANOVA with group (sleep or wake) and session (T1, T2) as between-subject variables (see figure 3.4 a). There was a main effect of session, with participants performing better at T1 (62.71% correct) compared to T2 (52.96% correct) as shown in figure 3.4a ($F[1,38] = 62.32$, $p < .001$). There was no interaction between group and session ($F[1,38] = .66$, $p = .42$). These results suggest overnight sleep was not beneficial in protecting memory of specific videos more than remaining awake. Rather, participants generally only got worse during the intervening time.

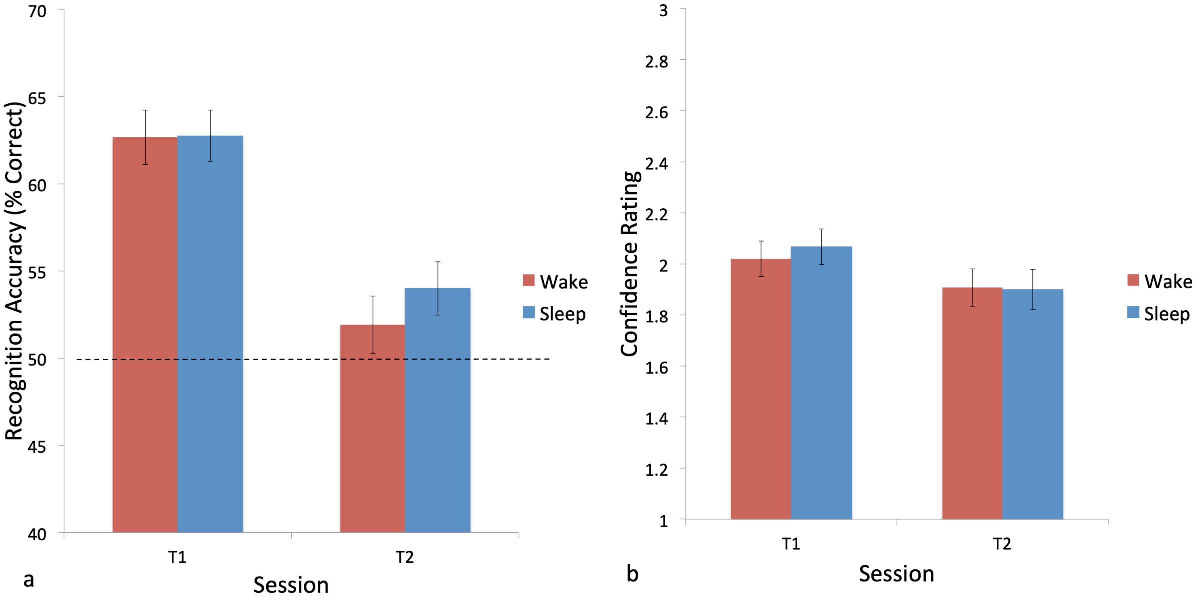

Figure 3.4 The graph on the left (a) represents participants' summed percent correct answers for both old and new videos between the two sessions (T1, T2) and between groups (wake, sleep). The black dashed line marks chance performance (50%). The graph on the right (b) represents participants' reported confidence between the two sessions (T1, T2) and between groups (wake, sleep). Higher values on the confidence scale represent higher confidence.

We also tested whether confidence ratings differed between the two groups using an ANOVA (see figure 3.4b), but only found a main effect of session ($F[1,38] = 8.52$, $p = .037$) and no interaction between the groups and sessions ($F[1,38] = .046$, $p = .83$). The results from the confidence measure indicate that participants' confidence reports generally matched their performance level.

To verify that our findings weren't a result of low-confidence trials increasing noise in our measure, we removed all lowest-confidence trials. The two-way repeated ANOVA showed a main effect of session ($p < .001$), but no interaction between groups and sessions ($p = .45$). The results mirror our earlier analysis with all trials and so we can conclude that low confidence trials weren't driving the effects.

Because T1 specificity accuracy hovered around chance (50%), we additionally wanted to confirm that our effects weren't due to noise induced by non-learners. We used a median split on both the sleep and wake groups and conducted an additional ANOVA using the top-half scorers on the specificity test (mean = 67.83%). The results mirrored those of the prior test (see figure 3.4a) in that we found a significant effect of session ($p < .001$) but no interaction between groups and sessions ($p = .28$). Thus, we concluded the lack of differences between the groups wasn't due to non-learners.

Categorization. On average, participants correctly selected the category of a video with 73.98% accuracy (SD = 15.34%). This is comparable to the average performance of 73.03% accuracy (SD = 16.49%) on the last three blocks of the training. Participants also rated their confidence for categorizing videos at 1.29 (SD = .45) on average.

We then conducted a two-way repeated measure ANOVA with group (sleep or wake) and session (T1, T2) as between-subject variables. We found no difference in performance between the sleep and wake groups at T1 for categorization of both old or new videos into the three categories (mean nap = 74.91%; mean wake = 72.81%; $t[35] = 0.34$, $p = .73$). Furthermore, the groups were not significantly different in identifying Alpha ($p = .80$), Bravo ($p = .38$), or Charlie ($p = .77$). Confidence ratings of categorization were not significant between the groups at T1 ($p = .76$), indicating that the sleep and wake groups were equally matched before the delay.

To assess the change in categorization accuracy from T1 to T2, we used a two-way repeated measure ANOVA that included group (sleep or wake), session (T1, T2), and item type (old or new video). There was a main effect of session (see figure 3.5a), with participants performing better at T1 compared to T2 ($F[1,38] = 29.68$, $p < .001$) and a main effect of item ($F[1,38] = 5.76$, $p = .02$) with participants performing better on the categorization of old videos more than new. There was no interaction between group and session (*p* = .79). There was an interaction between session and item (*F*[1,38] = 10.40, *p* = .002) such that participants' categorization of new videos dropped across the delay while categorization of old videos stayed relatively stable. The three-way interaction between group, session, and item was nearly significant (*F*[1,38] = 3.26, *p* = .07). These results show that neither sleep nor wake was beneficial to protecting memory for categorization as a whole, but that old videos were better categorized than new, especially so after the delay.

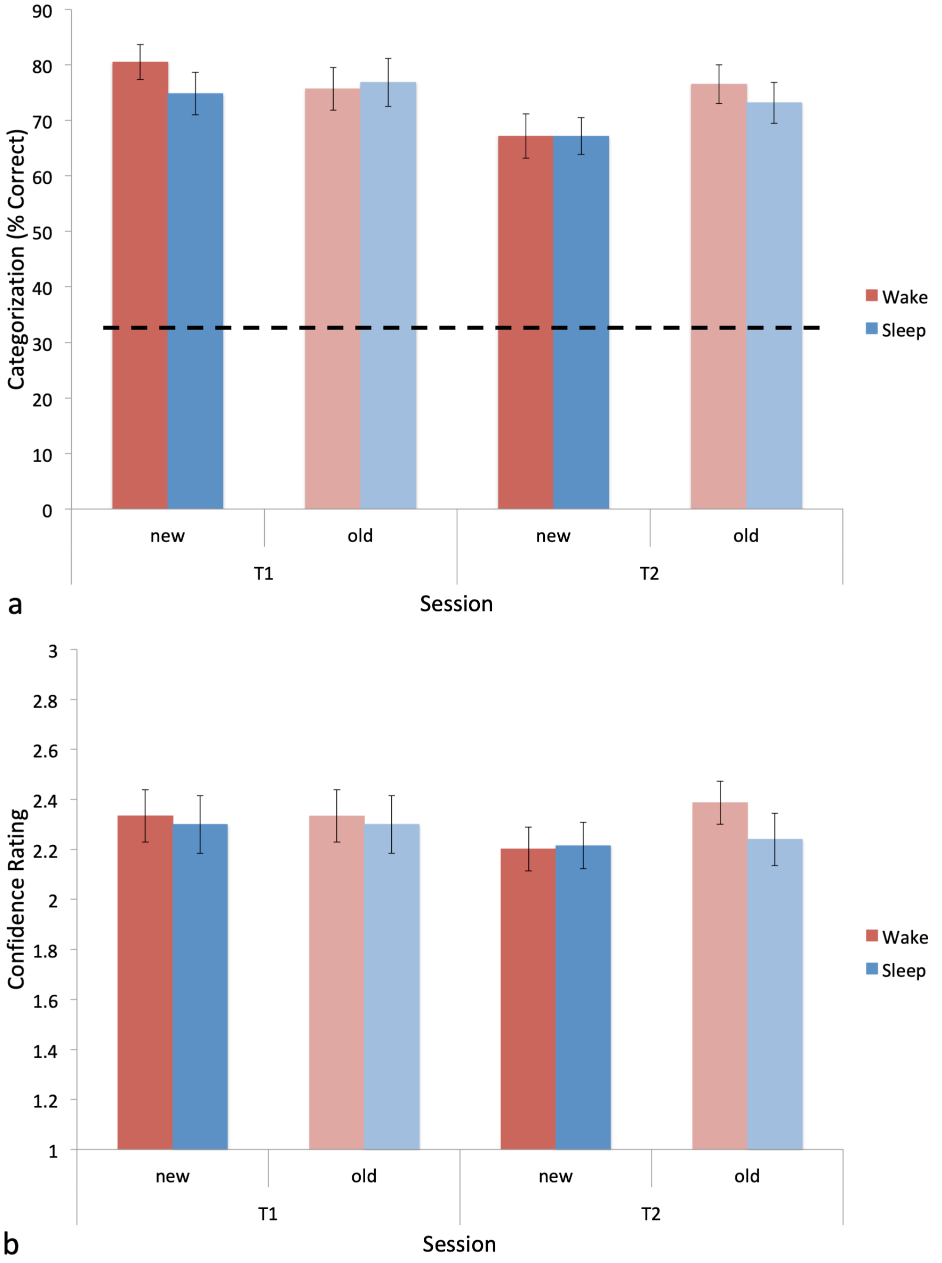

Figure 3.5 The top graph (a) represents participants' summed percent correct answers for categorization between the two sessions (T1, T2) and between groups (wake, sleep). Items are additionally split between old and new videos and marked with either bold (new) or pastel (old) coloring. The black dashed line marks chance performance (33.33%). The bottom graph (b) represents participants' reported confidence between the two sessions (T1, T2) and between groups (wake, sleep) for categorization. Higher values on the confidence scale represent higher confidence.

We also tested whether confidence ratings differed between the two groups across sessions and item types using an ANOVA. We did not find a main effect of session ($p = .14$), but did find a main effect of item type ($F[1,38] = 14.11$, $p < .001$) such that participants felt more confident on old rather than new videos. There were two significant 2-way interactions. First, the interaction between group and item type ($F[1,38] = 8.19$, $p = .006$) was significant, reflecting that the wake group rated their confidence about the same as the sleep group on new videos but were more confident on old videos, as shown in Figure 3.5b. Second, the interaction between session and item type ($F[1,38] = 14.11$, $p = .0006$) was significant, meaning participants felt less confident on new items at T2 but maintained their confidence on old items across the delay. Finally, there was a significant three-way interaction between group, session, and item ($F[1,38] = 8.19$, $p = .007$).

Summary of Study 1. The overall picture from Study 1 is after 9 blocks of training, participants were sufficiently able to distinguish both between old and new videos as well as recognize the videos' categories. We found that after either nighttime sleep or a period of wake over the day, both groups showed forgetting to barely above chance. There was also no significant interaction between the groups on categorization, though participants on average became worse at categorizing new videos after the delay and maintained categorization performance on old videos.

Participants' pattern of ratings of their confidence generally matched their actual performance. There were two exceptions to this, both on categorization. First, confidence on categorization generally didn't drop from T1 to T2, regardless of group. This is interesting as confidence for memory specificity did drop across time, as did actual performance. This may reflect that the training was solely based on categorization where participants received consistent feedback and thus felt more confident on the categorization tests. Secondly, there was an interaction such that the wake group rated their confidence higher for old items than the sleep group. This was not expected and doesn't have an obvious explanation.

Because performance on the specificity task essentially dropped to chance after the 12-hour delay, we hypothesized that a shorter delay would help protect against such catastrophic forgetting. Additionally, we were concerned that circadian factors could be influencing the results. Although the groups were matched for circadian preference and alertness at training, after the 12-hour delay they showed differences in alertness (as measured with the SSS) between the groups. Because of these concerns, we next ran an afternoon nap versus afternoon awake variant of this paradigm with EEG, allowing us to explore whether specific sleep features affect memory performance in this task. Additionally, this secondary design matched the time of day between groups, eliminating any potential circadian effects.

STUDY 2 RESULTS AND DISCUSSION

Sleepiness survey and circadian preference. The SSS scores did not differ between the nap and wake conditions for T1 (mean nap = 3.22; mean wake = 3.26; $t[35] = -0.12$, $p = 0.91$) or T2 (mean nap = ; mean wake = ; $t[24] = 1.76$, $p = 0.09$). Although it should be noted that 11 participants didn't complete a second SSS questionnaire (5 from the nap group and 6 from the wake). This suggests that there were no alertness differences between the groups, as would be expected because all participants began the study within the same time window. Additionally, the MEQ scores did not differ between the two groups (mean nap = 48.06; mean wake = 43.26; $t[35] = 1.81$, $p = 0.07$), suggesting that there were no differences in circadian preference between the two.

Training. In the second study, we were able to address the computer lag that wasn't recording button presses. Collapsing across the last three blocks, average performance across all participants was 67.60% (SD = 18.20%). As shown in figure 3.6, we found there were no differences in learning between the two groups on any block. Most importantly the two groups weren't significantly different at the final block (mean nap = 70.55%; mean wake = 68.77%; $t[35]$ =, $p = 0.75$), meaning their learning was equally matched.

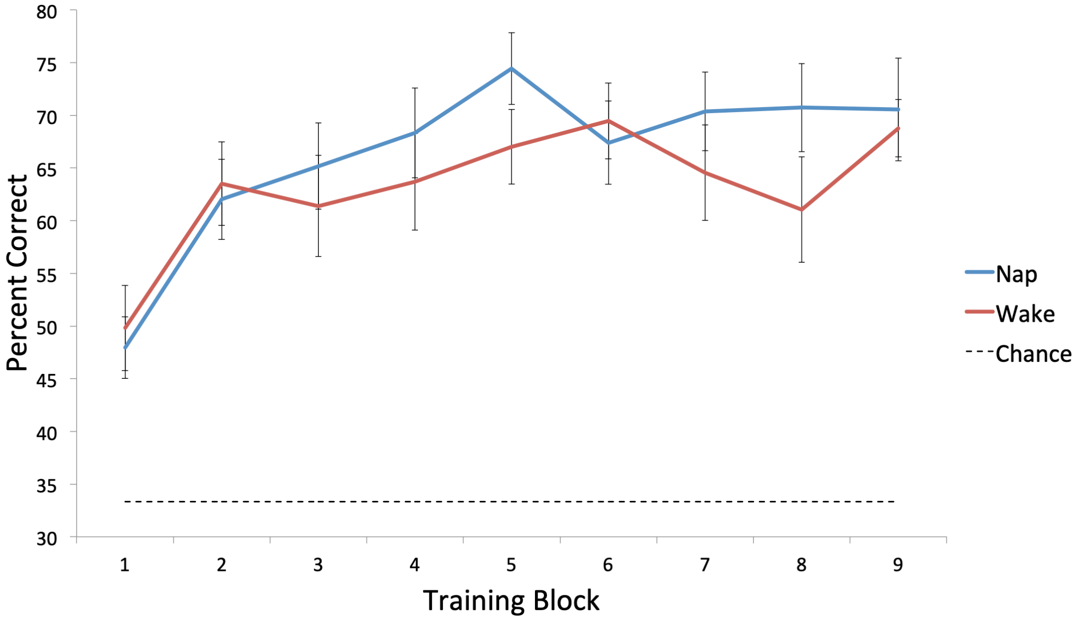

Figure 3.6. Line graph of average learning curves across the 9 training blocks for both the nap and wake groups separately. The x-axis denotes the training block and the y-axis denotes the average percent correct for that block. The dashed line represents chance (33.33%).

Memory Specificity. On average, participants correctly selected whether the video was old or new with 57.07% accuracy (SD = 6.21%), which is nearly identical to performance on study 1 (57.31% accuracy). On average, participants' confidence for specificity in this study was 2.20 (SD = .28). Examining just items called new, participants on average were correct (hit rate) for 52.61% of the trials and false alarmed (called a old video new) on 30.18% of the trials, resulting in a d-prime of .61 at T1 (z-scored hit rate = .07, z-scored false alarm = -.54). After the delay, the average hit rate was 49.28% (z-scored value = -.02) and average false alarm of 41.53% (z-scored value = -.23), resulting in a d-prime of .21 at T2. Although participants were randomly placed in each group, we did find a difference in performance between the nap and wake groups at T1 for recognition of whether a video was old or new (mean sleep = 62.96%; mean wake = 58.51%; *t*[35] = 2.07, *p* = .045). We further broke this test down by examining only old or new items and found the nap and wake groups weren't significantly different at recognizing only new videos (*p* = .35) or only old (*p* = .44). These T1 differences were unexpected as all participants began the experiment around the same time and were randomized into the two groups. Finally, we analyzed participants' confidence at T1 and found no difference between the groups (*p* = .97).

To assess the change in recognition accuracy from T1 to T2, we used a two-way repeated measure ANOVA with group (nap or wake) and session (T1, T2) as between-subject variables. As shown in figure 3.7, there was a main effect of session with participants performing better at T1 compared to T2 (*F*[1,35] = 36.79, *p* < .001). There was nearly a significant interaction between the nap and wake groups and session (*F*[1,35] = 3.04, *p* = .089). While these results might seem to indicate that the nap group saw an improvement over the delay compared to wake, this interaction isn't reliable and thus we cannot make any conclusive statements.

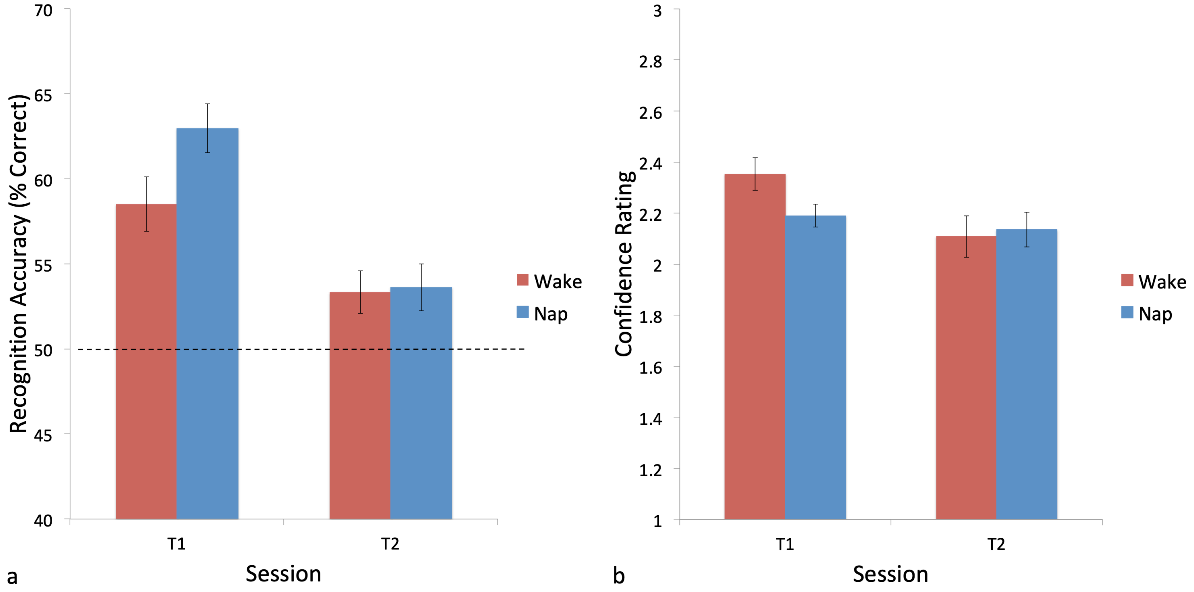

Figure 3.7 The graph on the left (a) represents participants' summed percent correct answers for both old and new videos between the two sessions (T1, T2) and identifies groups by color (red = wake, blue = nap). The black dashed line marks chance performance (50%). The graph on the right (b) represents participants' reported confidence between the two sessions (T1, T2) and between groups (wake, sleep).

We also tested whether confidence ratings differed between the two groups using an ANOVA, but only found a main effect of session ($F[1,35] = 14.20$, $p = .0006$) and no interaction between the groups and sessions ($F[1,35] = 1.04$, $p = .31$). Although the nap group was better at identifying old or new videos during T1, this did not translate to increased confidence for that test session. Confidence measures only seemed to track the decline in performance across the delay.

Similar to study 1, we wanted to verify that our findings weren't a result of low-confidence trials increasing noise in our measure. Hence, we removed all lowest-confidence trials (where confidence was rated as 1 for specificity). The two-way repeated ANOVA showed a main effect of session ($p < .001$), but no significant interaction between groups and sessions ($p = .06$). The results mirror our earlier analysis with all trials and so we can conclude that low confidence trials weren't driving the effects.

Lastly, we noted that T1 specificity accuracy was lower on this experiment compared to study 1 so we wanted to confirm that our effects weren't due to noise induced by non-learners. We used a median split on both the sleep and wake groups and conducted an additional ANOVA using the top-half scorers on the specificity test (mean = 65.42%). The results mirrored those of the prior test (see figure 3.7a), in that we found a significant effect of session ($p < .001$) but no interaction between groups and sessions ($p = .64$). Thus, we concluded the lack of differences between the groups wasn't due to non-learners.

Categorization. There were no differences in performance between nap and wake groups during T1 (mean nap = 75.83%; mean wake = 78.08%; $t[38] = 0.44$, $p = .65$). Furthermore, the groups were not significantly different in identifying Alpha ($p = .14$), Bravo ($p = .66$), or Charlie ($p = .39$). Unlike the specificity test, confidence ratings of categorization at T1 were not significantly different between the groups ($p = .31$), indicating that the nap and wake groups were matched before the delay.

Once again, we used a two-way repeated-measures ANOVA to assess the change in categorization accuracy from T1 to T2. This ANOVA included three factors: group (nap or wake), session (T1, T2), and item type (old or new video). There was a main effect of session, with participants performing better at T1 compared to T2 ($F[1,35] = 324.28$, $p = .01$) and a main effect of item ($F[1,35] = 11.86$, $p = .001$) with participants performing better on the categorization of old videos more than new. There was no interaction between the nap and wake groups and session ($p = .69$). There was an interaction between session and item ($F[1,35] = 13.90$, $p = .0006$) such that participants' categorization of new videos dropped across the delay while categorization of old videos stayed relatively stable. According to these results, neither nap nor wake was beneficial to protecting memory for categorization as a whole. Rather, old videos were better categorized than new videos and this was especially noticeable after the delay.

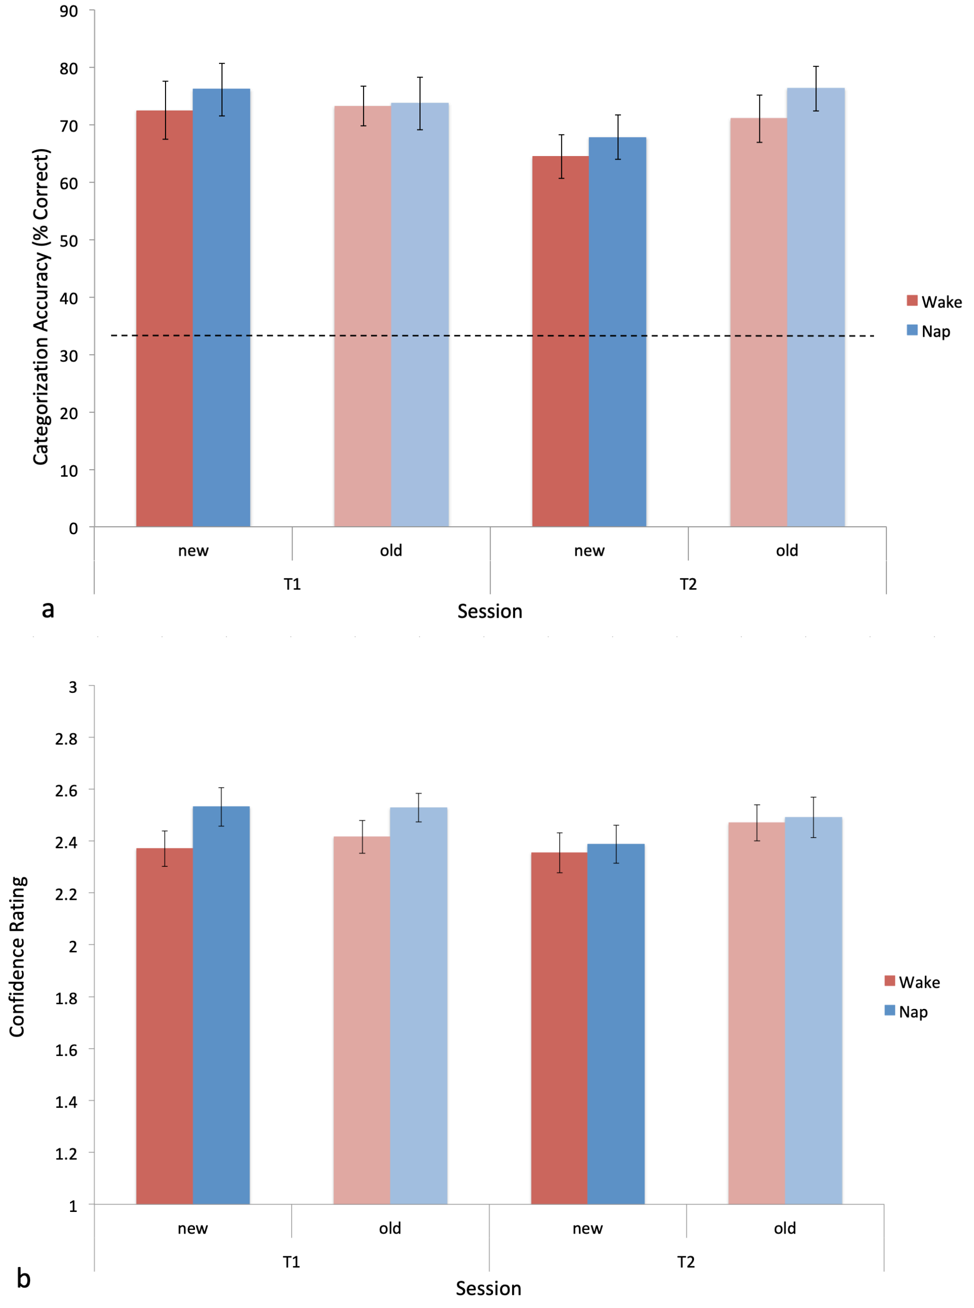

Figure 3.8 Top graph (a) represents participants' summed percent correct answers for categorization between the two sessions (T1, T2) and between groups (red = wake, blue = sleep). Items are additionally split between old and new videos and marked with either bold (new) or pastel (old) coloring. The black

dashed line marks chance performance (33.33%). The bottom graph (b) represents participants' reported confidence between the two sessions (T1, T2) and between groups (wake, sleep) for categorization.

Finally, we used an ANOVA with identical factors to those above (nap v. wake, T1 v. T2, old v. new stimuli) to examine confidence measures on categorization. The results showed a different pattern than the actual performance of participants on categorization. Like actual performance there was a main effect of item ($F[1,35] = 11.89$, $p = .001$) and there was an interaction between session and item ($F[1,35] = 7.81$, $p = .008$) such that participants' confidence rating of their categorization for new videos dropped across the delay while categorization for old videos stayed relatively stable. However confidence differed from performance in that there was no main effect of session ($p = .28$). These results show that while confidence ratings generally match actual performance, participants did not feel that their performance dropped between T1 and T2 when it actually did.

Electrophysiological Results. To test whether specific aspects of sleep influenced performance within the nap group (n = 18), we ran correlation analyses looking at time spent in sleep stages as well as sleep spindle factors (i.e. density and total number of fast and slow spindles). Our behavioral measures of interest were change in specificity memory (T2-T1), change in categorization memory, and change in confidence ratings for both specificity and generalization. Although we had expected to see a relationship between these sleep measures and behavior based on prior work, none of these measures significantly correlated with any of the sleep features.

	Wake (min)	N1 (min)	N2 (min)	SWS (min)	REM (min)	Average spindle density	Fast spindle density	Slow spindle density	Slow oscillation density	Total delta power (during NREM)
Average	26.76	14.474	27.68	21.08	2.45	4.55	2.52	2.02	4.09	85.36
SD	±12.44	±6.63	±8.74	±16.63	±4.78	±0.77	±0.74	±0.68	±3.74	±69.85

Table 3.1 Moving from left to right, the first 5 full columns in this table represent the average number of minutes spent in sleep stages (and standard deviation). The next 3 columns represent the average spindle

densities (number of spindles per minute in N2 or SWS) that are also divided between fast (13.5-15 Hz) and slow (12-13.5 Hz) spindles. Slow oscillation density in column 9 is defined as the number of 0.5-1 Hz oscillations per minute during SWS. Finally, the last column represents delta power (0.5-2 Hz) during N1, N2 and SWS.

We then reviewed other potential correlates between sleep and memory measures in an exploratory manner. Although these relationships were statistically significant on first glance, these correlations did not remain significant after correcting for multiple comparisons. We found a correlation between T2 specificity scores and slow spindle density ($r = 0.459$, $p = 0.048$) and all spindle densities ($r = 0.469$, $p = 0.043$), meaning nap participants who had more spindles (and slow spindles in particular) during sleep were better at recognizing old and new videos at T2. Relatedly, we found that T2 confidence ratings on the specificity task negatively correlated with minutes of SWS ($r = -0.542$, $p = 0.016$) meaning that participants who spent less time in SWS were more likely to rate their confidence more highly for the specificity task.

Summary of Study 2. Similar to study 1, participants learned across the 9 training blocks and were significantly above chance on categorization by the end of training. During T1, participants on average were above chance for categorization and recognition on the old/new video test.

On the memory specificity (i.e., recognition) test, the nap and wake groups performed significantly different at T1 with the wake group doing better. This was unexpected as participants were randomized into the two groups. We tried analyzing just the top-half performers, which did even out the T1 performance, but we found no interaction between groups across the delay. Similar to study 1, confidence measures on memory specificity mirrored actual performance.

The ANOVA results of the categorization task were matched those seen in study 1. Accuracy for this task dropped across the two sessions, though performance on new videos showed a greater decline than categorization of old videos. Once again, reported confidence matched these findings.

Finally, we found no sleep physiology signatures that significantly correlated with change in performance or confidence on either of the tests. We only found two correlations of note regarding memory specificity. Namely, participants who had higher spindle density, specifically of slow spindles, had better recognition memory at T2. We also found that participants who spent more time in SWS were less confident of their recognition memory at T2.

GENERAL DISCUSSION

We were interested in testing two major questions. First, whether a period of sleep as compared to wakefulness protected or improved memory for detailed recognition, categorization, or both. And secondly, whether forgetting or memory decay across time might impact these two forms of memory. In study 1, we showed that memory for specifics declines across a 12-hour delay, regardless of whether that period included wake or nighttime sleep. Additionally, we saw categorization declined over the delay regardless of group although to a lesser extent than recognition. We did find that participants were worse at categorizing the new videos compared to old, studied videos. To remove potential circadian confounds and reduce forgetting with a shorter delay we ran study 2. We once again showed forgetting on both recognition and categorization tests, regardless of group. And once again, forgetting was less catastrophic for categorization than memory for specific clips. In this section we consider the implications of these results for our understanding of memory changes across delays and whether highly confusable stimuli are the best choice for tackling these questions.

Dissociation between categorization and recognition indicates use of prototype-based memory. Although we saw a decrease in memory performance across the board, not all forms of memory were equally impacted. In both studies, we found that recognition memory dropped more than categorization performance. This is rather surprising as we exposed participants to the old videos 9 times each, which should have strengthened their memory for those particular exemplars. This finding supports the claim that participants were relying on a prototype memory representation to make decisions for categorization.

If participants were relying on specific examples they saw during learning, then their categorization should have decreased in a matching manner to what was observed on the old-new recognition test. Therefore, recognition of specific video clips was likely not driving their categorization performance old or new clips. This finding aligns with a previous study where participants had repeated exposure to training stimuli that should have benefitted exemplar-based decision-making, but actually supported the prototype theory (Minda & Smith, 2002).

There are a few caveats that should be recognized, but don't invalidate this argument. One potential push-back on this assertion is that high performance on prototype distortion tasks can be due to perceptual memory rather than a generalized prototype (Casale & Ashby, 2008; Little & Thulborn, 2006). However, perceptual tuning of the visual cortex would require static images like the dot paradigm and the stimuli used here were short video clips and so this cannot be a case of perceptual memory. Another concern is that the T1 performance was not matched between categorization and recognition for either study. Across the two studies, recognition of specific videos was an average of 61.72% while categorization was 75.41% correct at T1, and this was probably partially due to the training method used. Participants received feedback on categorization during the learning portion, but never got feedback on recognition memory. That said, across groups (sleep and wake) and studies (1 and 2), categorization only dropped an average of 4.97% while recognition of specific videos decreased 8.51%. Since categorization at T1 was generally quite good, we should expect participants to have further to drop but this wasn't the case. Thus, we conclude that this explanation doesn't fully account for differences in recognition and categorization. Instead, the available evidence indicates that forgetting of specific clips (i.e., exemplars) occurred within two hours but categorization performance decayed much more slowly and this pattern of results supports the prototype theory of categorization.

Categorization performance differs between old and new videos. In neither experiment did we observe differences between sleep and wake groups in categorization. We had originally hypothesized, in line with other theories, sleep would allow for transformation of memory such that categorization should improve after sleep (Inostroza & Born, 2013). However, we found that categorization performance dropped (to a lesser extent than recognition memory) after the delay, whether the delay included sleep or wake. This result is best explained by the amount of experience participants had with old videos compared to new. Participants saw each old video nine times across learning and practiced recognizing which category that video belonged to. Hence, it isn't surprising that participants were better able to categorize old videos.

No clear physiological markers for memory during sleep. Based on prior work, we were interested in two hypotheses regarding sleep and categorization in study 2. The first, based on several prior studies that looked at how information is reorganized during sleep (Beijamini et al., 2014; Monaghan et al., 2015; Verleger et al., 2013; Wagner et al., 2004), was that minutes of SWS would correlate with an improved performance on categorization. We did not find any evidence of this being the case. The second hypothesis was that minutes of REM sleep would correlate with improved categorization, similar to a previous study using the remote associates task (Cai et al., 2009). We did not find any evidence that minutes spent in REM correlated with any behavioral changes in either recognition or categorization.

We did find that recognition performance at T2 was positively correlated with slow spindle density. This is tempered by the fact that we didn't see a relationship between change in recognition (T2-T1) and any physiological markers from table 3.1. We can't be certain if slow spindles are a physiological marker for maintaining specificity or if people who generally have more slow spindles during sleep also tend to better differentiate between old and new videos. We are faced with a similar problem in interpreting the negative correlation between confidence at T2 for recognition memory and minutes of SWS. We once again found no relationship between change in confidence and SWS or any other sleep stages. Thus, we can't draw any strong conclusions regarding these two correlations.

Conclusion. Overall, our results suggest that sleep does not have a different effect on highly confusable video stimuli compared with wake. While previous research has found differences between sleep and wake over both short and 12-hour long delays (Graveline & Wamsley, 2017b; Hanert et al., 2017; Lau et al., 2011), a more recently published study also failed to find a difference between sleep and wake (Bailes et al., 2020; Tucker et al., 2020). Future studies should carefully select stimuli and training programs such that performance is well above chance on all tests. Training to a high level of performance will avoid interpretation problems, to prevent performance from declining to chance.

We also did not find a difference in categorization between sleep and wake, although there was a difference in performance between old and new videos. While participants were engaged with the visually rich content of the videos, future studies might want to use more differentiable stimuli.

The pattern of results between the categorization and recognition tests supported the prototype theory of memory categorization. Participants, on average, forgot specific videos more quickly than the decrease in identifying the correct category. Because they did not recall specific instances, they could not be relying on exemplars to determine categorization and thus were most likely relying on an abstracted prototype. We believe our results, which while they do add to the sleep and memory literature, are most impactful in distinguishing between the exemplar and prototype theories for forgetting across time.

CHAPTER 4

Experiment 2: Does memory reactivation affect specificity and generalization?[1]

Memory serves many functions in our daily lives. Two of those are recalling specific events and generalizing across many instances to form an abstract knowledge. For example, a scuba diver may recall a very specific memory of visiting the Great Barrier Reef and spotting a highly venomous stonefish resting on the reef's floor. This would be an example a specific, highly detailed memory. On the other hand, that same diver may visit the Solomon Islands, recognize a similar fish, and know that it's a midget stonefish, smaller but equally venomous. Memory researchers have long understood that these types of memories exist in parallel. Endel Tulving categorized them as semantic and episodic memory in his conceptualization of memory (Tulving, 1972). However, it's still an open question as to whether episodic and semantic memory are in opposition.

The complementary-learning-systems approach tackles this question by integrating the two forms of memories through different brain circuits (McClelland & O'Reilly, 1995; O'Reilly & Norman, 2002; Schapiro, Turk-Browne, et al., 2017). In the CLS model, the hippocampus rapidly encodes unique, individual memories and maintains the separation between the memories via pattern separation through sparse encoding (Bakker et al., 2008; Yassa & Stark, 2011). Then, over longer periods of time the cortex extracts shared features and generates a generalized, abstract transformed memory. Several researchers have argued that this process of extracting

---

[1] I would like to recognize my collaborators in this study, as they will be listed as co-authors once this chapter is submitted for publication. Sharon Noh, Ali Preston, and Victoria Lee have put in significant work in assisting with designing the stimuli, running participants, and providing critical advice all throughout the process.

regularities and generating a generalized concept can occur throughout the day, but may especially be prevalent during sleep (Ellenbogen, Hulbert, et al., 2006; Norman et al., 2005).

The process of stabilizing and integrating memories is referred to as memory consolidation. In the animal literature, memory consolidation is measured through the reactivation of firing patterns from cells that were involved those particular memories (Pavlides & Winson, 1989; Wilson & McNaughton, 1994). In human research it can be difficult to record neural activity at the cellular level and as such, most work in human memory consolidation relies in measuring behavioral and EEG correlates (Paller et al., 2020). Research into memory consolidation has focused heavily on sleep, namely because sleep seems to have the unique characteristic of lessened interference from external sources. Depending on the type of memory being studied, researchers have found that sleep, as compared to wake, is beneficial for memory. For example, sleep can protect specific memories (Hanert et al., 2017; Marshall & Born, 2007), and enhance generalization across memories to create schemas (Beijamini et al., 2014; Inostroza & Born, 2013). In some cases where researchers reported both specific and generalized memory, participants showed a trade-off between the two such that memory for specifics declined while generalization improved (Lau et al., 2011).

Rather than relying on studies that compare periods of sleep and wake, and thus contend with circadian and group differences, manipulating sleep through targeted memory reactivation (TMR) avoids such problems. Targeted memory reactivation is an experimental tool that allows researchers to direct endogenous memory reactivation and impact consolidation. During learning, participants are exposed to an auditory or olfactory cue simultaneously to the presentation of information they are meant to learn, such as the location of an item on a grid. That cue is then presented again during sleep, carefully to prevent rousing the participant to wakefulness, and it reactivates the memory such that the participant has better recall for the cued items compared to uncued.

Some of the earliest studies that established common methodology for using TMR focused on spatial memory (for review see Oudiette & Paller, 2013; Rasch et al., 2007; Rudoy et al., 2009). However, a recent meta-analysis of TMR studies found that spatial memory is one of many types of memory affected by sleep cueing (Hu et al., 2020). While the meta-analysis revealed that TMR has a consistent effect on spatial memory, other types of memory are affected as well. One way to view the spatial TMR findings is that they represent detailed, specific memory because higher accuracy requires detailed recall of the exact location.

Targeted memory reactivation can also be used to study generalization memory effects. For example, researchers in a 2016 study instructed participants to learn associations between ambiguous pictures and either positively or negatively valenced words before overnight sleep with half of the words cued. They found that cueing increased memory for recall of both the positive and negative learned associations such that participants remembered which word went with each picture. Cuing impacted how participants interpreted new picture-word pairings. If the participant was cued with negative words, they interpreted new negative words as less of a fit for the new image to which it was paired, but this pattern was reversed for new positive word-image pairings (Groch et al., 2016). This study indicates that sleep with cueing might also affect generalization.

Lexical competition, where a newly learned word competes with existing vocabulary, is another experimental design that hints at generalization. In one such experiment using lexical competition between words and non-words, the researchers hypothesized that TMR would lead to better lexical integration of the studied non-words. They did not find a direct effect of TMR on participants' response choices. However, they found that cueing when combined with a period of REM sleep was correlated with better integration of new information within existing knowledge, as measured by a slower reaction time for cued words compared to uncued words (Tamminen et al., 2017). In other words, participants were faster for uncued words and slower for cued words because the cued words were more confusable and they took longer to make a judgement. The authors argued that activation of words cued during SWS were then integrated during REM sleep such that participants took longer to respond to new non-words that shared phonetic attributes with the cued words. Together, these two studies hint that TMR may affect memory for generalized constructs as well as specifics. However, to date there have been no TMR experiments directly comparing cuing of generalization and specific memories.

In the present study, we used a set of landscape paintings by six relatively unknown artists to test both types of memory such that we could examine the influence of memory reactivation on them both. We intended to evaluate whether cueing an artist led to improved memory specificity on previously seen paintings, improved generalization memory for an artist's style in new paintings at the cost of specificity, or improved both. Participants studied a set of paintings and were tested both on memory for specific paintings and on generalized knowledge of each artists' style with respect to new paintings. We administered TMR during the nap and cued half of the artists and then again tested participants after wake. To ensure that participants were cued during slow-wave sleep (SWS), we recorded electroencephalography (EEG) throughout the nap.

There are several advantages of examining memory for paintings. First, the paintings could be used to test memory for specifics in a straightforward way. We arranged a multiple-choice recognition test with foils that were very similar to paintings seen during the learning phase of the experiment. Second, the paintings could also be used to test generalization memory, following the procedures used by Kornell and colleagues (2008) as stimuli because several of the selected artists have been reliably used to demonstrate generalization across an artists' style (Kornell et al., 2010; Kornell & Bjork, 2008). In the learning phase, participants gradually learned to identify the artist who had painted a set of paintings, such that they could identify the paintings of six different artists. Participants then used that knowledge in the generalization test when they attempted to identify the artist who had painted a set of new paintings. Additionally, the artists and paintings were selected to minimize the influence of prior knowledge, such that performance reflected information acquired during the learning phase. We selected four artists from the 2008 experiment (Judy Hawkins, Philip Juras, Ron Schlorff, George Wexler) and added an additional two (Richard McKinley & Jamie Grossman) from a more recent studies (Noh et al., 2020).

We hypothesized that generalization and memory specificity wouldn't be affected by TMR in the same way. Rather, we hypothesized that there would be a trade-off between the two memory types (Witkowski et al., 2020). We expected that memory consolidation, spurred on by TMR, would result in a more cortically-based memory that would protect generalization but harm memory specificity. Alternatively, though less expected, we thought that specificity might be protected such that generalization would be impaired.

Sleep Physiology

We were interested in several aspects of sleep physiology that have been previously associated with later memory performance. At the most basic level, the amount of time spent in various sleep stages has been associated with better memory, whether that be SWS or REM sleep

(Cairney, Durrant, Hulleman, & Lewis, 2014; Cairney, Lindsay, Sobczak, Paller, & Gaskell, 2016; Tamminen et al., 2017). One experiment found that rather than total time spent in a given stage, the number of times participants transitioned between REM and non-REM correlated with improved explicit awareness of sequences on a serial reaction time task (Kirov et al., 2015). The result demonstrates that SWS alone isn't enough to recognize a sequence hidden in the serial reaction time task, and that REM might play an important role. The importance of REM as well as non-REM in the 2015 study closely mirrors a 2017 experiment that found curing during SWS that was followed by REM sleep showed a cueing effect (Batterink, Westerberg, & Paller, 2017). Thus, we might expect that minutes spent in REM sleep lead to improved generalization while more SWS will improve recognition of old items (memory specificity), with cueing improving both only if REM sleep occurs over the course of the nap.

Moving past classic sleep stage scoring, researchers have found analyzing specific EEG frequencies explains more of the changes seen memory performance over sleep. Several experiments have found increased theta band power is related to improved memory performance after sleep with cueing (Antony, Cheng, et al., 2018; Farthouat et al., 2017; Schreiner & Rasch, 2015). Unfortunately, these studies didn't differentiate between memory specificity and generalization so the question of whether theta is a general indicator of improved memory or only relates to specific kinds of memory is still an open question. For example, increased theta power after a sleep cue could improve memory specificity, but not generalization. We analyzed both time spent in various sleep stages and theta power to see whether they correlated with behavior.

METHOD

Participants

We enrolled a total of 36 participants who were undergraduate students at Northwestern University or local Evanston community members. Seven participants were excluded from the final analyses; two did not sufficiently learn the material, one participant reported having insomnia, one reporting hearing the sound cues during the nap, and three were not cued because they never reached SWS. All included participants reported no history of neurological or sleep disorders. The final 29 participants (21 female, 1 non-binary, 7 male; 1 left-handed; mean age 21.07 yrs) were requested to wake up 2 hours earlier than their usual wake time and to abstain from caffeine on the day of the experiment. The Northwestern University Institutional Review Board approved this procedure and informed consent was obtained from all participants and everyone was paid for their participation.

Stimuli

The stimuli consisted of 108 paintings by 6 different artists (18 paintings each). Four of these artists were from the study by Kornell and Bjork (2008) and two were from the study by Noh et al. (2020). The paintings depicted landscapes or skyscapes that were cropped to remove the artists' names (i.e., McKinley, Schlorff, Hawkins, Wexler, Juras, Grossman). Using paintings by these relatively unknown artists avoided the potential complication of prior knowledge. All but one participant were naive to the artists' names and all were naive to the selected paintings. For the learning and recognition portions, 72 paintings (12 by each artist) were divided vertically into 3 equal-sized portions, called 'slices' here.

In addition, we used six auditory stimuli, each lasting 500 ms (hoot, meow, thunder, applause, cricket, bell). Each sound was assigned to an artist, without any clear relationships between the sounds, artists' names, or paintings.

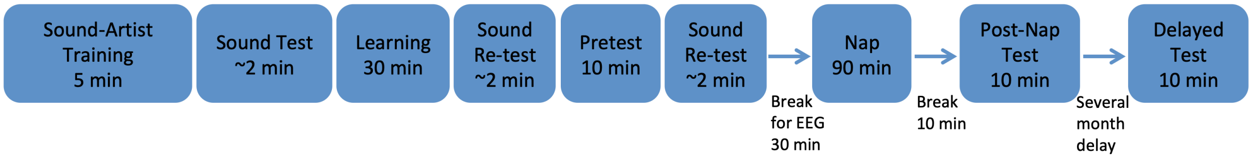

Figure 4.1 Flowchart representing the steps of the experiment and approximate amounts of time spent on each step.

Procedure

Participants arrived at the lab between noon and 2pm, consented to participate, and completed a Stanford Sleepiness Scale. They then began the eight-part study.

Sound-Name Association Training. During this phase of the experiment, participants were instructed to learn the sound-name pairings for a later test. Participants heard each sound cue, followed by a 500 ms delay before they saw the associated artist's name on the screen. The name remained on the screen for an additional 500 ms before the "next" button appeared on the screen and participants could advance to the next pairing. The list of six pairings was presented three times, each time in a different randomized order.

Sound-Name Association Test. One of the six sound cues was played through the speakers and participants attempted to select the paired artist name from a list of all the names. After selecting their answer, the participant was shown the correct name before moving to the next trial. The list of six pairings was tested twice and there was no time constraint on their selection.

Learning. On each learning trial, a painting slice appeared in the center of the screen (see Figure 4.2a). Participants had 10 s to select the artist for the painting slice from a list of the six artist names. After the participant responded, they were shown the correct name and the associated sound was played before advancing to the next painting. There were 30 trials within each block. Each block included six paintings, each by a different artist. The set of paintings were displayed five times, each time in a different random order. There were a total of six blocks. Paintings weren't repeated across blocks, such that participants studied a total of six paintings by each artist.

Sound-Name Association Re-test. To verify that participants learned the sound-name associations, they were tested again. The procedure was the same, except the list of six pairings was tested once.

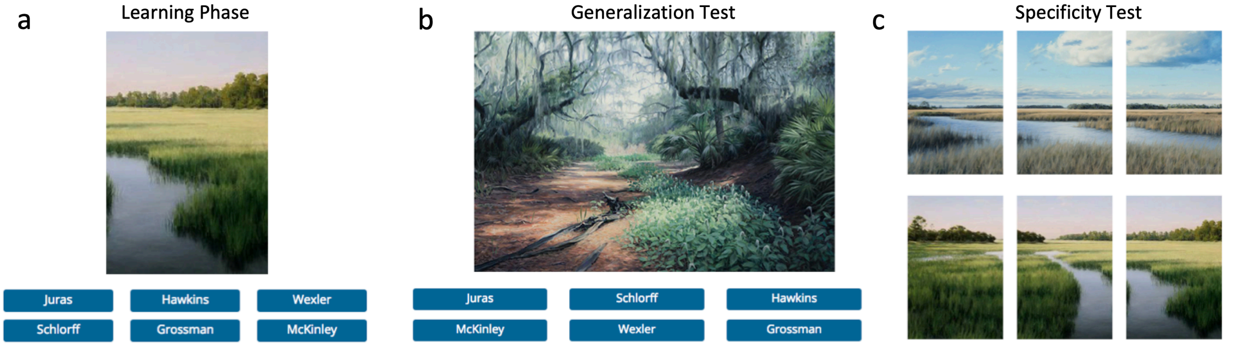

Figure 4.2 Examples of the three main portions of the task. Part a exemplifies the learning phase where participants saw a slice of the original painting and had to respond with the correct artist's name. Part b shows the generalization test where participants were shown a new, full painting and had to select the artist who painted it. Part c shows the specificity test where participants saw six slices (three from each painting) and had to select the exact slice they viewed during learning.

Pre-sleep Memory Tests. In order to test both memory specificity and abstraction, this phase consisted of two parts (see Figure 4.2bc). The generalization test included 18 paintings (3 by each artist) that were never seen during the learning phase. Each new painting was displayed individually per trial at the center of the screen, in a random order. Participants were instructed to select the artist they thought created the painting. They had 10 s to respond and were not given feedback.

In the memory specificity test, participants were tested on the painting slices from the learning phase. On each trial, participants saw an array of six slices in two rows. Each row contained three slices, ordered from left to right such that participants could see the full original paintings. The same artist created both paintings, which were similar, but one had never been seen before (lure). Participants attempted to select the slice from the learning phase within 10 s and weren't given feedback after their selection. There were 18 trials during the specificity test, 3 trials per artist, shown in a random order.

EEG Procedure. EEG was recorded using a Neuroscan EEG system, with a sampling rate of 250 Hz and a bandpass of 0.1 – 100 Hz. Electrodes in a cap were located at 21 scalp locations from the 10-20 system (Cz, C3, C4, Fpz, Fp1, Fp2, Fz, F3, F4, F7, F8, Pz, P3, P4, T3, T4, T5, T6, Oz, O1, O2, and two mastoids). Additionally, two electrodes were placed on the face, one lateral to the right eye and one below the left eye, for the electro-oculogram. One electrode was placed on the chin for the electromyogram. Impedance was lowered below 5 kΩ for all electrodes. EEG preprocessing was completed using the EEGlab package (Delorme & Makeig, 2004) for Matlab (The Mathworks, Munich, Germany). Data were re-referenced to the averaged mastoids and any noisy scalp channels were interpolated using the spherical interpolation method in EEGlab.

Nap. The participant's futon was converted from a chair to a bed position and covered with sheets, a pillow, and a blanket. White noise quietly played through the speakers (around 35-40 dB) and participants were told they would have a 90-min nap opportunity. Sleep physiology was monitored online and when the participant reached SWS for at least two 30-s epochs, half of the sound cues were presented in random order. Each sound cue was 500 ms in duration with a 3500-ms interstimulus interval between the end of the previous cue and the start of the next.

Sounds were played continuously while the participant was in SWS. If there was an arousal during SWS, cuing was paused and did not resume until at least two 30-s epochs of SWS occured. After the 90 min ended, participants were awakened if they weren't already awake. If the participant was still in SWS, the experimenter waited up to 15 min so that the participant could be woken from lighter sleep.

Post-sleep Memory Tests. After waking, participants were given a short break to get water or use the restroom before beginning the post-sleep tests. The tests were identical to the pre-sleep tests, but with different stimuli. Participants were tested on 18 new paintings for generalization that weren't seen during learning or the pre-sleep test. The memory specificity test was composed of the 18 painting slices from learning that they weren't tested on pre-sleep, paired with 18 never-before-seen painting slices by the matching artist. Additionally, participants completed the sound-name association test a final time. All participants correctly identified at least five of the six artists.

Debriefing. Participants were asked whether they heard any disturbances during the nap period. If they reported hearing a sound, they were asked to specify if it was one of the sound cues or another disruption.

Analyses

Behavioral. The main measure in this study was the percent change between pre- and post-sleep test scores. To examine the effect of TMR on memory, we conducted two-way ANOVAs with factors cuing (Cued, Uncued) and time (Pre-Sleep, Post-Sleep). Separate ANOVAs were conducted for the generalization and memory specificity tests. All tests were two-tailed (alpha = 0.05).

Sleep staging. Naps were scored offline using the sleep SMG package in Matlab (http://sleepsmg.sourceforge.net). The 90-min nap was scored in 30-s epochs using scalp electrodes (C3, Fz, Fp1, P4, Oz) and labeled as wake, stage 1, stage 2, SWS, or REM. In order to ensure accurate sleep scoring, the rater had no knowledge of when sounds were presented. Each sleep stage was then correlated with behavior and the TMR effect.

RESULTS

Behavioral

Sleepiness survey and circadian preference. Participants reported an average sleepiness rating of 3.21 (SD = 0.83) on the Stanford Sleepiness Scale, where the scale ranges from 1-7 and smaller values indicate greater wakefulness. Participants also completed the Morningness-Eveningness Questionnaire after the post-sleep test and reported an average score of 42.82 (SD = 10.03), indicating neither a preference for the morning nor evening.

Sound-name association accuracy. Participants were highly accurate at identifying the sound cues that matched each artist’s name. Because there were 6 alternatives, chance was 16.67%. On the first sound-name association test, accuracy was 91.09% (SD = 28.52%). On the second sound-name association test, accuracy was 96.55% (SD = 18.30%). On the final test, accuracy remained at 96.55% (SD = 18.30%). By the last test, 26 out of the 29 participants were 100% accurate for recognizing the sound-artist association.

Generalization performance. Participants correctly identified the artists of novel paintings with 57.85% accuracy (SD = 16.42%, range = 22.22% - 94.44%) during the pre-sleep memory test. Again, because there were 6 alternatives, chance was 16.67%.

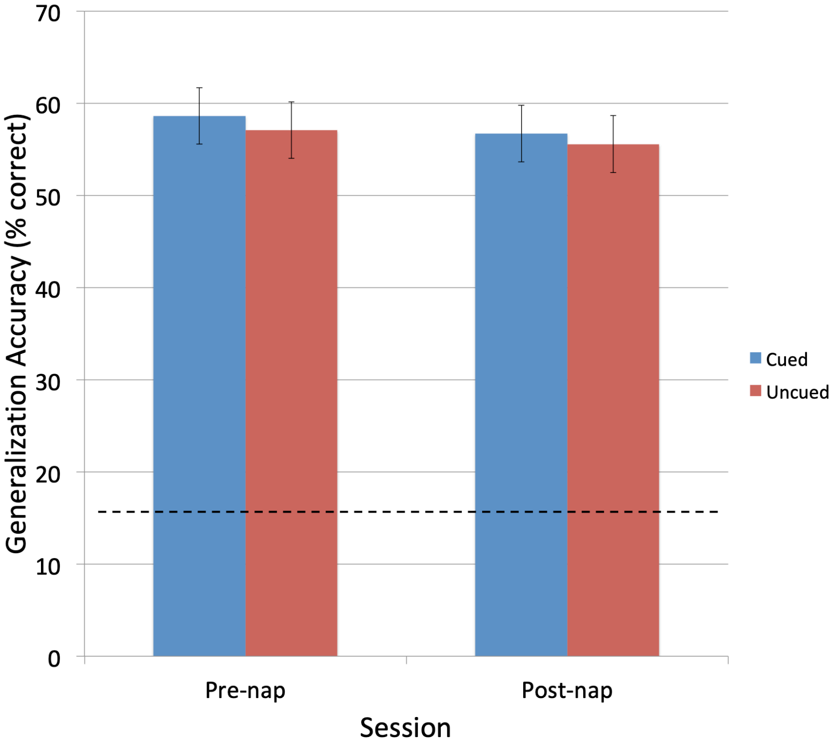

Figure 4.3 Graph represents average percent correct performance on the generalization test for cued (blue) and uncued (red) artists for both before (pre-nap) and after (post-nap) TMR. The dashed line represents chance performance (16.67%).

Next, we assessed the change in generalization accuracy from before to after sleep and the effect of TMR. We found no difference between the scores before and after sleep ($F[1, 27] = 0.25$, $p = 0.62$). We found no difference between the cued and uncued conditions ($F[1, 27] = 0.42$, $p = 0.52$), nor did it interact with sessions ($F[1, 27] = 0.01$, $p = 0.93$). These results indicate that cuing during sleep did not affect participants' ability to correctly identify the artist of a new painting.

Specificity performance. On the specificity test, participants correctly identified the artist during the pre-sleep memory test with 60.73% accuracy (SD = 13.90%, range = 38.89% - 88.89%). As in the generalization test, chance performance was 16.67%. If participants could rule out three of the choices by remembering the correct painting but not the correct slice, they would perform at 33.33% accuracy. Yet, pre-sleep accuracy was always higher.

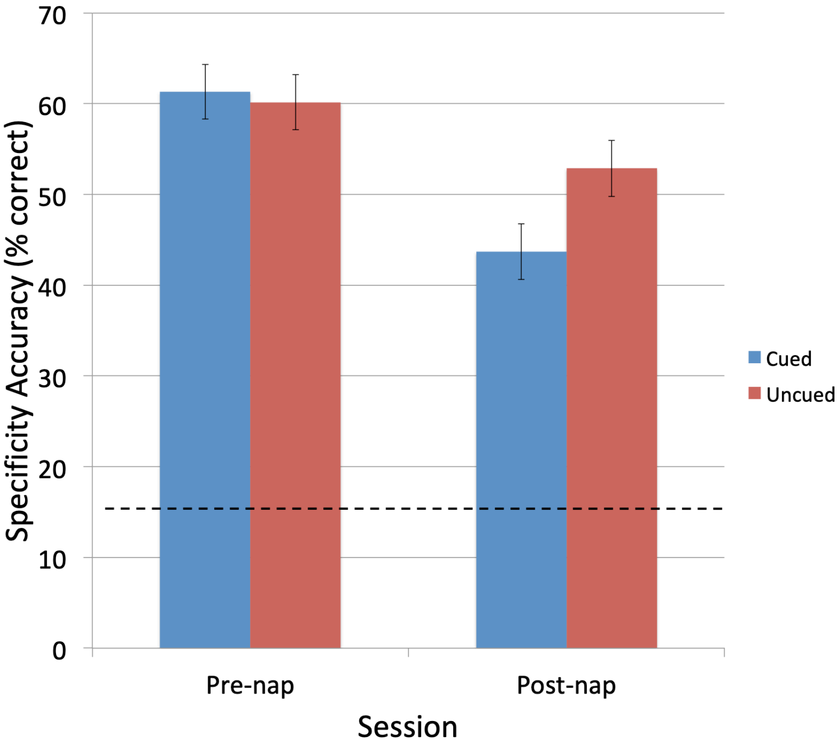

Figure 4.4 This graph represents average percent correct on the memory specificity test for cued (blue) and uncued (red) artists for both before (pre-nap) and after (post-nap) TMR. The dashed line represents chance performance (16.67%).

We then assessed the change in specificity accuracy from before to after sleep as well as the effect of TMR. Accuracy declined after sleep ($F[1, 27] = 32.55$, $p < 0.001$). We found no difference between the cued and uncued conditions ($F[1, 27] = 3.68$, $p = 0.07$). There was an interaction effect such that performance declined to a different degree for cued and uncued conditions ($F[1, 27] = 5.77$, $p = 0.02$). Follow-up testing showed no difference between cued and uncued conditions before the nap ($t[28] = 0.31$, $p = 0.75$), but there was a difference between the conditions after the nap ($t[28] = -2.18$, $p = 0.03$). Thus, cuing produced worse performance for the cued artists compared to the uncued artists.

Reaction Times. We also included an analysis to determine if TMR influenced speed of response. Overall, response times averaged 4.07 s (SD = 1.01) on the generalization test and 5.39 s (SD = 1.17) on the specificity test.

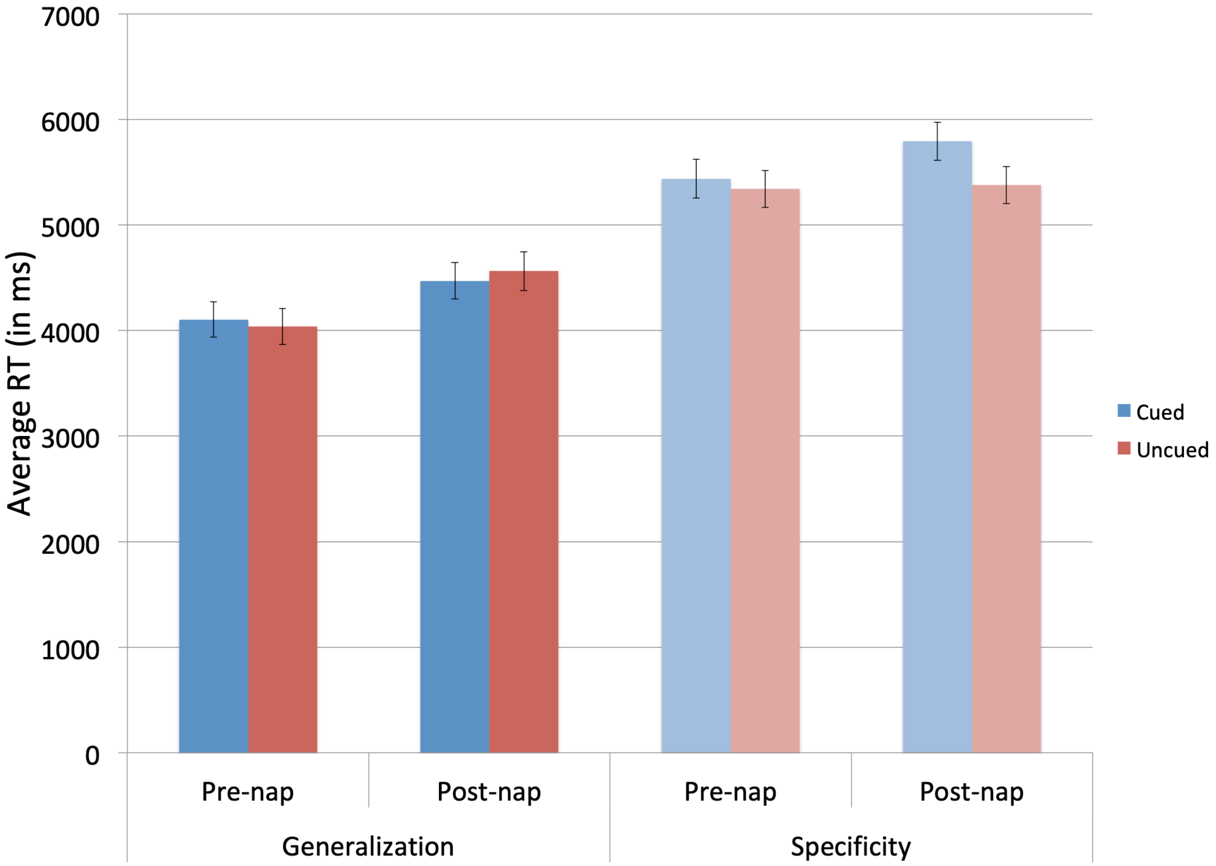

Figure 4.5 This graph represents the average reaction times for both tests (dark colors are generalization, light colors are specificity). The cued artists are depicted in shades of blue while the uncued artists are depicted in shades of red. Note that while analyses were conducted after RTs were z-scored by participant, this graph depicts raw values for ease of interpretability.

Reaction times were *z*-scored by each participant to reduce individual variation for analyses. For generalization, responses were faster before the nap than afterwards (*F*[1, 27] = 9.74, $p < 0.005$). For specificity, responses were faster in the cued condition than the uncued condition (*F*[1, 27] = 4.34, $p = 0.046$). Neither the generalization nor the specificity test reaction times showed a significant interaction between session and cuing. Notably, assignment of artists to cued and uncued conditions was made to balance accuracy but not reaction time across conditions. Given that reaction time differences were apparent in the specificity test before sleep, they could not have been produced by the cueing manipulation.

Physiological Correlates of Behavioral Change

Theta Power. Previous research using TMR has found a relationship between theta power and memory in that increased theta power during cued sleep was positively correlated with later memory specificity (Farthouat et al., 2017; Schreiner & Rasch, 2015). In this experiment, we did not see an increase in theta power after cuing, and thus those results are not reported here.

	Wake (min)	N1 (min)	N2 (min)	SWS (min)	REM (min)
Average	21.66	15.78	29.6	21.21	3.86
SD	±14.25	±10.56	±11.38	±11.31	±5.00

Table 4.1 This table includes the average number of minutes (and standard deviation) spent in each stage of sleep by participants.

Sleep Stages. We also hypothesized that time spent in various sleep stages would relate to later memory performance. For example, based on one study where participants who reached REM during their nap displayed increased memory abstraction, we expected minutes of REM to correlate with performance on the generalization task (Tamminen et al., 2017). While minutes in REM did not correlate with improvement for participant's generalization accuracy, time spent in REM was negatively correlated with the reaction time of the generalization cuing benefit (p = 0.03, r = -0.39). Cuing benefit was calculated as the subtraction of cued artists (post nap minus pre) minus uncued artists. This correlation shows that more time spent in REM sleep correlates with a larger improvement on the cuing benefit of reaction time for generalization. However, only 15 participants entered REM sleep during the study. When we repeated this analysis excluding participants who did not get REM sleep, the correlation was no longer significant (p = 0.21).

Results from Testing after a Longer Delay

Participants were contacted several months after their initial visit and asked if they were willing to return for a third memory test (mean = 105.71 days, SD = 28 days). This memory test was identical to the previous tests, except that it differed in the number and type of stimuli. Participants were tested on 17 new paintings for generalization that weren't seen ever before. Because we were unable to find enough stimuli by one artist (Ron Schlorff), we re-used one generalization painting from before the nap to give us a full set of 18 generalization paintings.

The memory specificity test included 36 pairs of paintings displayed in an identical manner to the two previous tests. Because we had already tested the memory for all painting slices from the learning portion, we decided to re-test everyone's memory for all of the paintings rather than only select a few. Participants were not tested on the sound-artist associations for this test, but were asked to name the sounds associated with each artist as best they could from memory. Twenty-one participants (72% of original sample) returned to complete this additional memory test (13 female, 1 non-binary, 7 male; 1 left-handed; mean age of 20.76 yrs

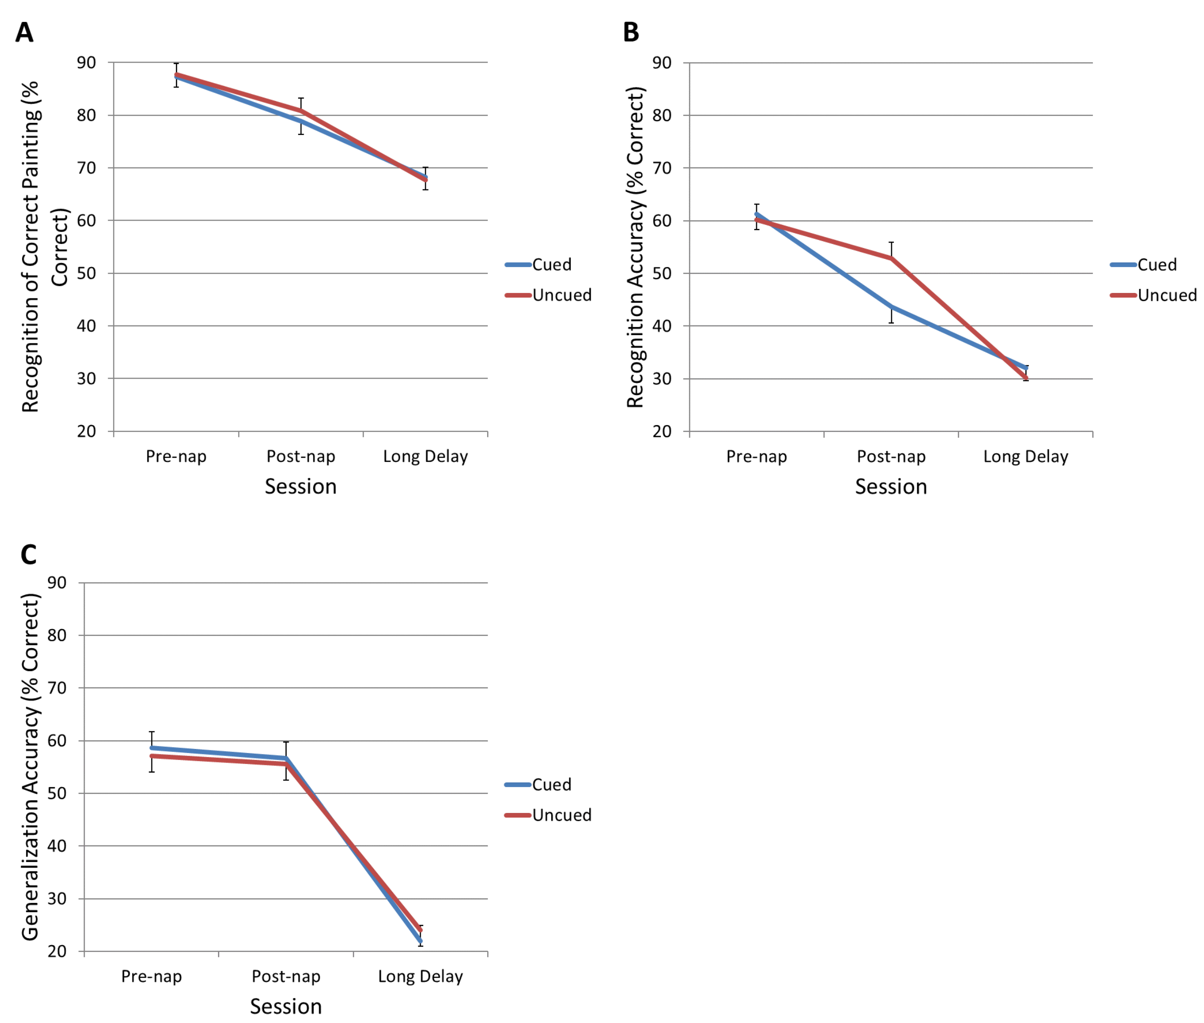

Figure 4.6 This figure is composed of three graphs. Graph A represents performance on the recognition (specificity) test for selecting the correct painting, but not the correct slice on the test. The x-axis spans across the three test points (before sleep, after sleep, several months later) for both cued (blue line) and uncued (red line) conditions. The y-axis represents the percent correct for a participant on average. The x- and y-axis and color-coding are identical across all three graphs. Graph B represents performance on the recognition (specificity) test for selecting the correct slice. Finally, graph C represents performance on the generalization test.

After this long of a delay, we did not find a significant difference between cued and uncued artists for generalization or specificity.

DISCUSSION

In this study, we sought to test targeted memory reactivation’s impact on generalization and memory specificity using paintings by relatively unknown artists. Results showed that while generalization of artists' styles was not affected by targeted memory reactivation, memory for specific paintings was made worse by cuing. We then conducted a follow-up test several months later (average of 105.71 days) to see whether this pattern of results changed. After the long delay, all cuing effects were abolished and accuracy on both memory measures dropped precipitously (31.14% accuracy for recognition, 21.91% accuracy for generalization).

We did not find a correlation among participants who had REM sleep between minutes of REM and a cuing benefit for either generalization or memory specificity, although this may be due to the low number of participants who reached REM sleep during the nap. Additionally, we did not find the amount of time spent in SWS or lighter sleep (N1 & N2) correlated with a change in performance.

We were surprised to find that TMR of artists led to poorer performance on the recognition task that served as our specificity measure without the expected commensurate improvement in generalization. However, we should note that previous studies using TMR on spatial tasks don't show a direct improvement on cued items (Cairney et al., 2016; Creery et al., 2015; Rudoy et al., 2009; van Dongen et al., 2012). Rather, prior research has found that cuing protects memories from forgetting and does not necessarily provide an additional boost above the pre-sleep baseline. Because generalization performance didn't decrease after the nap for either cued or uncued artists, it's difficult to know whether later forgetting would affect the uncued artists more than the cued ones. We did try to examine this question by conducting a follow-up test, but this was not initially part of our design. Participants returned an average of 105.71 days after learning and were able to recall very little from the prior session (approximately 4 paintings were correctly identified on average, 3 out of the 18 paintings would place a participant at chance performance). While it is tempting to argue a follow-up session a few days to a week later might've shown an effect of cueing, further research is needed to confirm whether this kind of generalization learning is indeed affected by TMR.

On the specificity recognition test, we did find a cuing effect such that accuracy of cued artists decreased after the nap. In the case of previous studies that failed to find a TMR effect, this was due to the timing of the cues being played too close together (Farthouat et al., 2017; Schreiner, Lehmann, et al., 2015). The study reported here avoided such a confound by playing each cue a full 1500ms after the end of the presentation of the previous cue, thus avoiding potential interference with ongoing memory processing.

So far, only a single published study has found TMR to negatively impact memory for the cued items (Göldi & Rasch, 2019). In an at-home application of TMR, the authors found that when TMR decreased memory accuracy when it woke participants or disrupted their sleep based on self-report. While we did not formally record how rested participants felt post-nap, we can look at the electrophysiological recordings to see whether cues disrupted participants sleep. Previous research in the lab has found that when TMR cues disrupt sleep, the number of minutes spent in SWS correlates with the cuing decrement (cued- uncued) (Whitmore et al., under review). We did not find a significant correlation between participant's minutes in SWS and cuing effect in this study ($p = 0.47$).

Additionally, we wanted to test if cuing specifically affected selecting the wrong slice or led participants to select a painting slice from the lure painting on the recognition test. We did not find an interaction between cuing and session ($F[1, 27] = 0.11$, $p < 0.74$), mean cuing did not make participants more likely to select a slice from the lure painting. Targeted memory reactivation only impacted memory at the 'slice-level' rather than between the two paintings.

One potential explanation could be that TMR led to an increase in boundary extension. When observers are asked to recall a picture or scene, they will often include more imagery than was presented and this is called boundary extension (Hubbard et al., 2010). For example, if you saw a photo of part of Chicago's skyline you would most likely report seeing more of the city. Boundary extension tends to occur more for scene-oriented images like landscapes as opposed to object-centered images (Bainbridge & Baker, 2020), and occurs even if participants are warned ahead of time (Intraub & Bodamer, 1993). This fits well with the design of our study. Participants viewed landscapes that were indeed truncated from the original paintings and while they were warned of later memory tests, this would not necessarily ablate boundary extension.

However, this potential explanation isn't totally satisfactory. Boundary extension occurs with milliseconds after exposure, and is generally regarded as a part of online integration (Dickinson & Intraub, 2008). No work has been done to test whether sleep affects boundary extension, let alone TMR. To test whether cuing led to boundary extension, we conducted an analysis on studied painting slices that were either at the far left or far right of the original painting (excluding the studied center slices). We then counted the number of slices away participants were from the studied image. We tested whether there was a difference between cued and uncued artists on this error measurement, with the expectation that if cuing led to boundary extension then those participants would likely only be a single slice removed from the true image. We did not find a difference between the errors in cued and uncued artists, but this could be because there were so few test items in this analysis.

Our preferred explanation for these findings is that there was a tradeoff between memory specificity and generalization due to TMR. The decrease in recognition memory could be an example of losing unnecessary specificity (ie., participants get the correct painting but not the correct slice) in order to help generalization. Unfortunately, TMR effects are generally only revealed when there is significant forgetting, which we didn't have for the generalization test. Future studies would benefit from a more appropriately timed follow-up test that would show some forgetting on the generalization task, without their performance falling to near chance. Additionally, this study was somewhat limited in the number of artists and paintings that were used which limited the strength of analyses like our examination of TMR-induced boundary extension. Future studies should consider using a greater number of stimuli to increase the power of their findings. Overall, this study expands our knowledge of memory consolidation through targeted memory reactivation. We conclude that TMR does not impact all forms of memory identically, and further work is needed to understand how cuing modifies memory consolidation during sleep.

CHAPTER 5

Future Directions and Conclusion

Summary

The dissertation studies on the whole contribute novel information to the scientific literature on sleep and memory. In these studies, I found memory for specific instances and generalization across multiple instances are forgotten across different time scales and are differently impacted by TMR. In chapter two, I lay out why TMR is a particularly powerful tool for examining consolidation with respect to memory specificity and generalization. Then in chapter three, I describe results from testing memory for videos of movement through naturalistic outdoor environments demonstrating that memory for categorization fades much more slowly than recognition memory. Finally, in chapter four I use sounds during sleep to bias memory consolidation and alter memory for specific paintings but not performance on generalization.

In the two studies included in chapter three, I used scenes that belonged to three categories to test both categorical knowledge as well as recognition of specific scenes (i.e., identified old or new) after a delay including sleep or wakefulness. Based on previous work, I expected that sleep would play a protective role and memory would be better after sleep compared to wake (Friedrich et al., 2020; Graveline & Wamsley, 2017a; Hanert et al., 2017). While I was surprised to find no difference between the sleep and wake groups, regardless of the delay, there was a significant difference between the two types of memory. This may have been due to the unique stimuli we used; namely, the videos showed moving though naturalistic outdoor environments that have never been used for a study like this before. Regardless of group, memory for specifics declined noticeably over the delay (either 12 hours or 2 hours), whereas categorization stayed relatively high. This pattern matches the pattern of results seen in chapter four. Even using a totally different stimuli set and a different training paradigm, there was still a significant drop on memory specificity but not generalization. It was only at the long delay of several months in chapter four that there was a significant drop in generalization. The findings in both chapters three and four provide additional support for the prior claim that generalization and memory for specifics rely on different memory mechanisms.

In fact, the TMR results of chapter four additionally support this assertion. After a nap with TMR of half the studied artists, only memory for specific paintings was affected. There was no difference in generalization performance between cued and uncued artists. On the specificity test, cuing artists led to participants performing worse on recognizing the specific paintings by those artists at test. This was quite unexpected as most TMR studies find that cuing improves memory for some specific measures like distance on the spatial task (Creery et al., 2015; Hu et al., 2020; Rudoy et al., 2009). Because there was not a significant decrease in performance on generalization, it's possible that the effect of TMR wasn't measurable for that test. Targeted memory reactivation usually manifests as a protection against forgetting rather than a direct increase in performance (Hu et al., 2020). If there isn't enough forgetting or enough variance in a test, perhaps TMR still operates on those stimuli but the effect isn't large enough to be measured. As a feature of the design added later or something, I attempted to collect data on participants' performance after a longer time period with the expectation that there would be enough forgetting in the generalization task to differentiate between cued and uncued artists. Unfortunately, collecting responses after an average of 105 days meant that we had too long a delay and performance dropped to barely above chance. In the next section, I will outline three additional follow-up experiments to better address the issues of: 1) time delay on memory, 2) how sleep physiology during TMR affects the cuing results, and 3) a follow-up experiment to answer whether the TMR-specificity effect can be explained by boundary extension.

Future Directions

Time delays. In chapter four, we attempted to bring back the majority of the participants for testing their memory at a delay. However, because data collection was initially completed in the fall, by the time the students (a majority of our sample) returned to campus it had been a minimum of 2 months since they initially participated. Perhaps it should then be unsurprising that participants struggled on both memory tests. Even using identical stimuli and design, it would have been much more appropriate to complete the follow-up test at a 2-week delay. Prior work using paintings by Sharon Noh (Noh and others, personal communication) has shown that associative memory can still be reliably measured two weeks after learning. Unfortunately, there are no published studies tracking the forgetting curve for memory specificity and generalization on matched tests. Ideally, a large set of additional participants could be recruited for the learning phase, with different subsets of the participants brought in at various time points from a day after learning up to a month after learning to plot the forgetting curve and make a data-driven decision for the timing of the follow-up test. However, such a study would require a heavy investment of time and resources that would exceed the value of such a pilot.

Spindle and slow-oscillation coupling. As previously discussed in this thesis, SWS often appears in correlational analyses as a significant predictor for later memory performance. During SWS (and stage 2), fast bursts of neural activity, called sleep spindles, often occur. Spindles are usually broken into slow spindles with a mean frequency of 11-13.5 Hz, and fast spindles with a mean frequency of 13.5-16 Hz (Mölle et al., 2011). In fact, sleep spindles and the cross-frequency coupling of those sleep spindles with slow oscillations during SWS have been previously linked with cued memory consolidation improvements (Batterink et al., 2016; Cairney et al., 2014; Cousins et al., 2014; Farthouat et al., 2017; Groch et al., 2017; Schreiner, Göldi, et al., 2015).

The phase of slow oscillations during sleep has influenced memory consolidation in prior studies (Batterink et al., 2016), but only recently have researchers focused on the endogenous frequency coupling of sleep spindles and slow oscillations (Canolty & Knight, 2010). Some researchers have argued that brain regions communicate with each other through the coordination of fast oscillations (like sleep spindles) by slow oscillations, thus creating unified representations of memory (Düzel et al., 2010). In this case, the slower oscillations serve as long range connections between brain regions while the faster oscillations represent local processing (Varela et al., 2001). Memory representations require many different brain regions to orchestrate the retrieval of any given memory and cross-frequency coupling provides a mechanism to achieve this need.

In addition to the analyses for chapter four, we could examine whether cross-frequency coupling between sleep spindles and slow oscillations predicts later memory performance. I am currently working with sleep researcher Daniela Grimaldi, MD, PhD to verify the sleep staging for participants. After all participants have been sleep scored a second time, we anticipate examining two aspects of their sleep. First, we will look at when cues were played during SWS and how that relates to the phase of the underlying oscillation. In other words, we will test whether the timing of cues in relation to the slow oscillations impacts the cuing effect. Secondly, we will examine the cross-frequency coupling of sleep spindles and slow oscillations during

SWS. Based on previous research, we expect that participants who have more consistent coupling between the frequencies will show a larger TMR effect after sleep than those who had less consistent coupling.

Boundary extension follow-up study. The results of chapter four hinted at an interesting possibility not previously addressed by the literature- whether TMR can affect boundary extension of visual stimuli. When cued during sleep, participants were worse at correctly identifying which painting slice by the cued artist they saw at learning. But cuing did not make participants more likely to select a slice from the lure painting, perhaps indicating that participants' errors were due to boundary extension. If cuing led to increased boundary extension, one would expect to see more errors towards neighboring slices and not an increase in errors made by selecting slices from the lure painting.

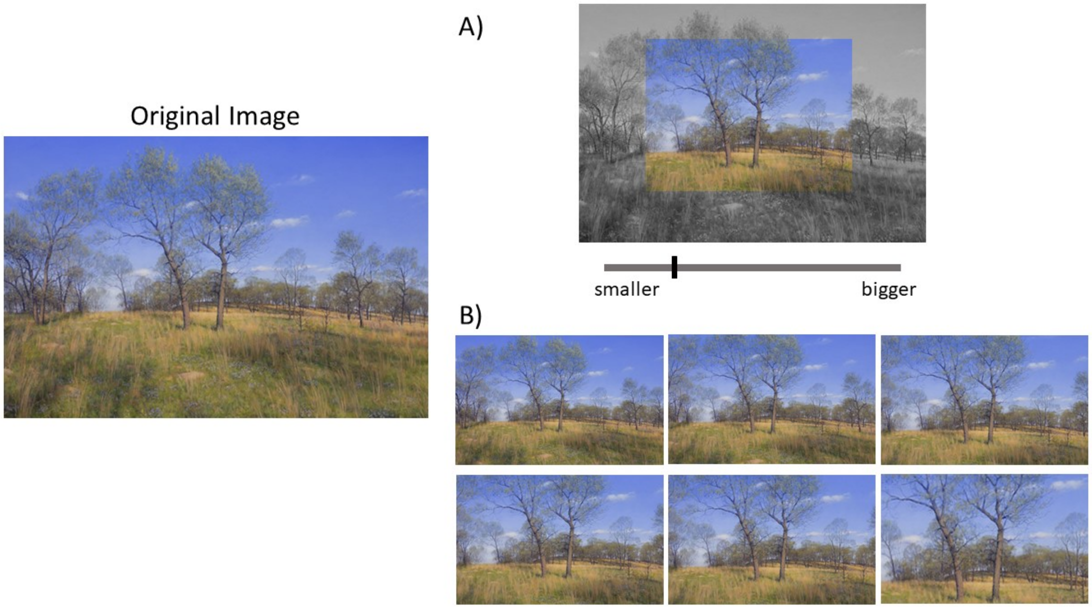

Figure 5.5 On the left is an example of a painting by Juras that could be used in a new experiment to examine boundary extension. Participants could be tested for their recognition of the boundaries either through a) a sliding scale where they try to match the image as precisely as possible, or b) select the studied image from an array of six choices.

One way to verify that this was the case would be to re-design the stimuli for the memory specificity test. Rather than having participants choose from multiple painting slices, stimuli could be generated by ‘zooming in’ at different levels from the initial painting and then presenting all images as the same size at test as shown in Figure 5.1. Boundary extension makes the claim that people report having seen more information in the peripheral areas of a scene than they actually saw (Hubbard et al., 2010). Thus, one might expect that in this new design, cuing would lead participants to select an image that includes more information than what they actually studied (i.e., a more zoomed out image). In order to maintain some balance between the generalization and specificity tests, six versions of each painting could be generated in this manner (ex., 17%, 32%, 49%, 64%, 81%, and 100% the original size) with one of the versions randomly selected for inclusion in the learning phase (shown in Figure 5.1a). Alternatively, if the focus is completely on boundary extension then it might be better to allow participants to respond using a sliding scale (shown in Figure 5.1b). Rather than limiting responses to a one-in-six forced choice, a sliding scale would allow for more variance and could be a stronger tool for addressing this unanswered question.

Conclusion

Despite a long history of investigating the differences between types of memory, there remain gaps in our knowledge with particular regard to how sleep affects consolidation of these different memory types. Researchers have established that sleep is important for consolidation of a wide array of memory types. However, little work has been done to determine how detail-specific memory and generalization might relate to each other antagonistically, and how that relationship might be affected by sleep consolidation. Through the experiments described in this dissertation, I have provided evidence that the time course of consolidation is different for generalization and memory specificity. I showed that the rate of forgetting differs between the two memory types, and targeted memory reactivation may not affect both in either the same way or to the same degree. My work has demonstrated the benefits and detriments of experimental design choices when investigating different memory types and provides a guide for future work on this topic.

REFERENCES

Alger, S. E., & Payne, J. D. (2016). The differential effects of emotional salience on direct associative and relational memory during a nap. *Cognitive, Affective, & Behavioral Neuroscience*. https://doi.org/10.3758/s13415-016-0460-1

Antony, J. W., Cheng, L. Y., Brooks, P. P., Paller, K. A., & Norman, K. A. (2018). Competitive learning modulates memory consolidation during sleep. *Neurobiology of Learning and Memory*, *155*, 216–230. https://doi.org/10.1016/j.nlm.2018.08.007

Antony, J. W., Piloto, L., Wang, M., Pacheco, P., Norman, K. A., & Paller, K. A. (2018). Sleep spindle refractoriness segregates periods of memory reactivation. *Current Biology*, *28*(11), R656–R677. https://doi.org/10.1016/j.cub.2018.04.020

Antony, J. W., Schönauer, M., Staresina, B. P., & Cairney, S. A. (2018). Sleep Spindles and Memory Reprocessing. *Trends in Neurosciences*. https://doi.org/10.1016/j.tins.2018.09.012

Ashby, F. G., & Maddox, W. T. (2011). Human category learning 2.0. *Annual Review of Psychology*, *56*(1), 147–161. https://doi.org/10.1146/annurev.psych.56.091103.070217

Bailes, C., Caldwell, M., Wamsley, E. J., & Tucker, M. A. (2020). Does sleep protect memories against interference? A failure to replicate. *PLoS ONE*, *15*(2), 14.

Bainbridge, W. A., & Baker, C. I. (2020). Boundaries extend and contract in scene memory depending on image properties. *Current Biology*, *30*(3), 537-543.e3. https://doi.org/10.1016/j.cub.2019.12.004

Bakker, A., Kirwan, C. B., Miller, M., & Stark, C. E. (2008). Pattern separation in the human hippocampal CA3 and dentate gyrus. *Science*, *319*(5870), 1640–1642.

Banks, W. P. (2000). Recognition and source memory as multivariate decision processes. *Psychological Science*, *11*(4), 267–273. https://doi.org/10.1111/1467-9280.00254

Batterink, L. J., Creery, J. D., & Paller, K. A. (2016). Phase of spontaneous slow oscillations during sleep influences memory-related processing of auditory cues. *Journal of Neuroscience*, *36*(4), 1401–1409. https://doi.org/10.1523/JNEUROSCI.3175-15.2016

Batterink, L. J., & Paller, K. A. (2017). Sleep-based memory processing facilitates grammatical generalization: Evidence from targeted memory reactivation. *Brain and Language*, *167*, 83–93. https://doi.org/10.1016/j.bandl.2015.09.003

Batterink, L. J., Westerberg, C. E., & Paller, K. A. (2017). Vocabulary learning benefits from REM after slow-wave sleep. *Neurobiology of Learning and Memory*, *144*, 102–113. https://doi.org/10.1016/j.nlm.2017.07.001

Beijamini, F., Pereira, S. I. R., Cini, F. A., & Louzada, F. M. (2014). After being challenged by a video game problem, sleep increases the chance to solve it. *PLoS ONE*, *9*(1), e84342. https://doi.org/10.1371/journal.pone.0084342

Bendor, D., & Wilson, M. A. (2012). Biasing the content of hippocampal replay during sleep. *Nature Neuroscience*, *15*(10), 1439–1444. https://doi.org/10.1038/nn.3203

Boyce, R., Williams, S., & Adamantidis, A. (2017). REM sleep and memory. *Current Opinion in Neurobiology*, *44*, 167–177. https://doi.org/10.1016/j.conb.2017.05.001

Busemeyer, J. R., Dewey, G. I., & Medin, D. L. (1984). Evaluation of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*(4), 638–648. https://doi.org/10.1037/0278-7393.10.4.638

Cai, D. J., Mednick, S., Harrison, E., Kanady, J., & Mednick, S. (2009). *REM, not incubation, improves creativity by priming associative networks*. *106*(25), 10130–10134.

Cairney, S. A., Ashton, J. E., Roshchupkina, A. A., & Sobczak, J. M. (2015). A dual role for sleep spindles in sleep-dependent memory consolidation? *Journal of Neuroscience*, *35*(36), 12328–12330. https://doi.org/10.1523/JNEUROSCI.2463-15.2015

Cairney, S. A., Durrant, S. J., Hulleman, J., & Lewis, P. A. (2014). Targeted memory reactivation during slow wave sleep facilitates emotional memory consolidation. *Sleep*, *37*(4), 701–707. https://doi.org/10.5665/sleep.3572

Cairney, S. A., Guttesen, A. á V., El Marj, N., & Staresina, B. P. (2018). Memory consolidation is linked to spindle-mediated information processing during sleep. *Current Biology*, *28*(6), 948-954.e4. https://doi.org/10.1016/j.cub.2018.01.087

Cairney, S. A., Lindsay, S., Sobczak, J. M., Paller, K. A., & Gaskell, M. G. (2016). The benefits of targeted memory reactivation for consolidation in sleep are contingent on memory accuracy and direct cue-memory associations. *Sleep*, *39*(5), 1139–1150. https://doi.org/10.5665/sleep.5772

Canolty, R. T., & Knight, R. T. (2010). The functional role of cross-frequency coupling. *Trends in Cognitive Sciences*, *14*(11), 506–515. https://doi.org/10.1016/j.tics.2010.09.001

Casale, M. B., & Ashby, F. G. (2008). A role for the perceptual representation memory system in category learning. *Perception & Psychophysics*, *70*(6), 983–999. https://doi.org/10.3758/PP.70.6.983

Cousins, J. N., El-Deredy, W., Parkes, L. M., Hennies, N., & Lewis, P. A. (2014). Cued memory reactivation during slow-wave sleep promotes explicit knowledge of a motor sequence.

Journal of Neuroscience, *34*(48), 15870–15876. https://doi.org/10.1523/JNEUROSCI.1011-14.2014

Cox, R., Hofman, W. F., & Talamini, L. M. (2012). Involvement of spindles in memory consolidation is slow wave sleep-specific. *Learning & Memory*, *19*(7), 264–267. https://doi.org/10.1101/lm.026252.112

Creery, J. D., Oudiette, D., Antony, J. W., & Paller, K. A. (2015). Targeted memory reactivation during sleep depends on prior learning. *SLEEP*. https://doi.org/10.5665/sleep.4670

Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21. https://doi.org/10.1016/j.jneumeth.2003.10.009

Dickinson, C. A., & Intraub, H. (2008). Transsaccadic representation of layout: What is the time course of boundary extension? *Journal of Experimental Psychology: Human Perception and Performance*, *34*(3), 543–555. https://doi.org/10.1037/0096-1523.34.3.543

Diekelmann, S., & Born, J. (2010). The memory function of sleep. *Nature Reviews Neuroscience*. https://doi.org/10.1038/nrn2762

Doxey, C. R., Hodges, C. B., Bodily, T. A., Muncy, N. M., & Kirwan, C. B. (2018). The effects of sleep on the neural correlates of pattern separation. *Hippocampus*, *28*(2), 108–120. https://doi.org/10.1002/hipo.22814

Düzel, E., Penny, W. D., & Burgess, N. (2010). Brain oscillations and memory. *Current Opinion in Neurobiology*, *20*(2), 143–149. https://doi.org/10.1016/j.conb.2010.01.004

Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D., & Walker, M. P. (2007). Human relational memory requires time and sleep. *Proceedings of the National Academy of Sciences*, *104*(18), 7723–7728.

Ellenbogen, J. M., Hulbert, J. C., Stickgold, R., Dinges, D. F., & Thompson-Schill, S. L. (2006). Interfering with theories of sleep and memory: Sleep, declarative memory, and associative interference. *Current Biology*, *16*(13), 1290–1294. https://doi.org/10.1016/j.cub.2006.05.024

Ellenbogen, J. M., Payne, J. D., & Stickgold, R. (2006). The role of sleep in declarative memory consolidation: Passive, permissive, active or none? *Current Opinion in Neurobiology*, *16*(6), 716–722. https://doi.org/10.1016/j.conb.2006.10.006

Farthouat, J., Gilson, M., & Peigneux, P. (2017). New evidence for the necessity of a silent plastic period during sleep for a memory benefit of targeted memory reactivation. *Sleep Spindles & Cortical Up States*, *1*(1), 14–26. https://doi.org/10.1556/2053.1.2016.002

Foster, D. J. (2017). Replay comes of age. *Annual Review of Neuroscience*, *40*(1), 581–602. https://doi.org/10.1146/annurev-neuro-072116-031538

Friedrich, M., Mölle, M., Friederici, A. D., & Born, J. (2020). Sleep-dependent memory consolidation in infants protects new episodic memories from existing semantic memories. *Nature Communications*, *11*(1), 1298. https://doi.org/10.1038/s41467-020-14850-8

Gais, S., & Born, J. (2004). Declarative memory consolidation: Mechanisms acting during human sleep. *Learning & Memory*, *11*(6), 679–685.

Glisky, E. L., Polster, M. R., & Routhieaux, B. C. (1995). Double dissociation between item and source memory. *Neuropsychology*, *9*(2), 229–235. https://doi.org/10.1037/0894-4105.9.2.229

Göldi, M., & Rasch, B. (2019). Effects of targeted memory reactivation during sleep at home depend on sleep disturbances and habituation. *Npj Science of Learning*, *4*(1), 5. https://doi.org/10.1038/s41539-019-0044-2

Gómez, R. L., Bootzin, R. R., & Nadel, L. (2006). Naps promote abstraction in language-learning infants. *Psychological Science*, *17*(8), 670–674.

Graveline, Y. M., & Wamsley, E. J. (2017a). The impact of sleep on novel concept learning. *Neurobiology of Learning and Memory*, *141*, 19–26.

Graveline, Y. M., & Wamsley, E. J. (2017b). The impact of sleep on novel concept learning. *Neurobiology of Learning and Memory*, *141*, 19–26.

Groch, S., McMakin, D., Guggenbühl, P., Rasch, B., Huber, R., & Wilhelm, I. (2016). Memory cueing during sleep modifies the interpretation of ambiguous scenes in adolescents and adults. *Developmental Cognitive Neuroscience*, *17*, 10–18. https://doi.org/10.1016/j.dcn.2015.10.006

Groch, S., Schreiner, T., Rasch, B., Huber, R., & Wilhelm, I. (2017). Prior knowledge is essential for the beneficial effect of targeted memory reactivation during sleep. *Scientific Reports*, *7*(1). https://doi.org/10.1038/srep39763

Hanert, A., Weber, F. D., Pedersen, A., Born, J., & Bartsch, T. (2017). Sleep in humans stabilizes pattern separation performance. *The Journal of Neuroscience*, *37*(50), 12238–12246. https://doi.org/10.1523/JNEUROSCI.1189-17.2017

Hasselmo, M. E. (1999). Neuromodulation: Acetylcholine and memory consolidation. *Trends in Cognitive Sciences*, *3*(9), 9.

Hauner, K. K., Howard, J. D., Zelano, C., & Gottfried, J. A. (2013). Stimulus-specific enhancement of fear extinction during slow-wave sleep. *Nature Neuroscience*, *16*(11), 1553–1555. https://doi.org/10.1038/nn.3527

Homa, D., Rhoads, D., & Chambliss, D. (1979). Evolution of conceptual structure. *Journal of Experimental Psychology*, *5*(1), 11–23.

Homa, D., Sterling, S., & Trepel, L. (1981). Limitations of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *7*(6), 418–439. https://doi.org/10.1037/0278-7393.10.4.638

Honma, M., Plass, J., Brang, D., Florczak, S. M., Grabowecky, M., & Paller, K. A. (2016). Sleeping on the rubber-hand illusion: Memory reactivation during sleep facilitates multisensory recalibration. *Neuroscience of Consciousness*, *2016*(1). https://doi.org/10.1093/nc/niw020

Hu, X., Cheng, L. Y., Chiu, M. H., & Paller, K. A. (2020). Promoting memory consolidation during sleep: A meta-analysis of targeted memory reactivation. *Psychological Bulletin*, *146*(3), 218–244. https://doi.org/10.1037/bul0000223

Hubbard, T. L., Hutchison, J. L., & Courtney, J. R. (2010). Boundary extension: Findings and theories. *Quarterly Journal of Experimental Psychology*, *63*(8), 1467–1494. https://doi.org/10.1080/17470210903511236

Inostroza, M., & Born, J. (2013). Sleep for preserving and transforming episodic memory. *Annual Review of Neuroscience*, *36*(1), 79–102. https://doi.org/10.1146/annurev-neuro-062012-170429

Intraub, H., & Bodamer, J. L. (1993). Boundary extension: Fundamental aspect of pictorial representation or encoding artifact? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*(6), 1387–1397.

Jegou, A., Schabus, M., Gosseries, O., Dahmen, B., Albouy, G., Desseilles, M., Sterpenich, V., Phillips, C., Maquet, P., Grova, C., & Dang-Vu, T. T. (2019). Cortical reactivations during sleep spindles following declarative learning. *NeuroImage*, *195*, 104–112.

Jenkins, J. G., & Dallenbach, K. M. (1924). Obliviscence during sleep and waking. *The American Journal of Psychology*, *35*(4), 605. https://doi.org/10.2307/1414040

Ji, D., & Wilson, M. A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience*, *10*(1), 100–107. https://doi.org/10.1038/nn1825

Kim, B., Kocsis, B., Hwang, E., Kim, Y., Strecker, R. E., McCarley, R. W., & Choi, J. H. (2017). Differential modulation of global and local neural oscillations in REM sleep by homeostatic sleep regulation. *Proceedings of the National Academy of Sciences*, 201615230. https://doi.org/10.1073/pnas.1615230114

Kirov, R., Kolev, V., Verleger, R., & Yordanova, J. (2015). Labile sleep promotes awareness of abstract knowledge in a serial reaction time task. *Frontiers in Psychology*, *6*. https://doi.org/10.3389/fpsyg.2015.01354

Klinzing, J. G., Kugler, S., Soekadar, S. R., Rasch, B., Born, J., & Diekelmann, S. (2018). Odor cueing during slow-wave sleep benefits memory independently of low cholinergic tone. *Psychopharmacology*, *235*(1), 291–299. https://doi.org/10.1007/s00213-017-4768-5

Kornell, N., & Bjork, R. A. (2008). Learning concepts and categories: Is spacing the “enemy of induction”? *Psychological Science*, *19*(6), 585–592.

Kornell, N., Castel, A. D., Eich, T. S., & Bjork, R. A. (2010). Spacing as the friend of both memory and induction in young and older adults. *Psychology and Aging*, *25*(2), 498–503. https://doi.org/10.1037/a0017807

Lau, H., Alger, S. E., & Fishbein, W. (2011). Relational memory: A daytime nap facilitates the abstraction of general concepts. *PLoS ONE*, *6*(11), e27139. https://doi.org/10.1371/journal.pone.0027139

Lerner, I., & Gluck, M. A. (2019). Sleep and the extraction of hidden regularities: A systematic review and the importance of temporal rules. *Sleep Medicine Reviews*, *47*, 39–50. https://doi.org/10.1016/j.smrv.2019.05.004

Lewis, P. A., & Durrant, S. J. (2011). Overlapping memory replay during sleep builds cognitive schemata. *Trends in Cognitive Sciences*, *15*(8), 343–351. https://doi.org/10.1016/j.tics.2011.06.004

Lewis, P. A., Knoblich, G., & Poe, G. (2018). How memory replay in sleep boosts creative problem-solving. *Trends in Cognitive Sciences*, *22*(6), 491–503. https://doi.org/10.1016/j.tics.2018.03.009

Little, D. M., & Thulborn, K. R. (2006). Prototype-distortion category learning: A two-phase learning process across a distributed network. *Brain and Cognition*, *60*(3), 233–243. https://doi.org/10.1016/j.bandc.2005.06.004

Lustenberger, C., Boyle, M. R., Alagapan, S., Mellin, J. M., Vaughn, B. V., & Fröhlich, F. (2016). Feedback-Controlled Transcranial Alternating Current Stimulation Reveals a Functional Role of Sleep Spindles in Motor Memory Consolidation. *Current Biology*, *26*(16), 2127–2136. https://doi.org/10.1016/j.cub.2016.06.044

Lutz, N. D., Diekelmann, S., Hinse-Stern, P., Born, J., & Rauss, K. (2017). Sleep supports the slow abstraction of gist from visual perceptual memories. *Scientific Reports*, *7*, 42950. https://doi.org/10.1038/srep42950

Maquet, P., Laureys, S., Peigneux, P., Fuchs, S., Petiau, C., Phillips, C., Aerts, J., Del Fiore, G., Degueldre, C., Meulemans, T., Luxen, A., Franck, G., Van Der Linden, M., Smith, C., & Cleeremans, A. (2000). Experience-dependent changes in cerebral activation during human REM sleep. *Nature Neuroscience*, *3*(8), 831–836. https://doi.org/10.1038/77744

Marshall, L., & Born, J. (2007). The contribution of sleep to hippocampus-dependent memory consolidation. *Trends in Cognitive Sciences*, *11*(10), 442–450. https://doi.org/10.1016/j.tics.2007.09.001

McClelland, J. L., & O’Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, *102*(3,419-457).

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

Medin, D. L., & Smith, E. E. (1984). Concepts and concept formation. *Annual Review of Psychology*, *35*(1), 113–138.

Mednick, S. C., McDevitt, E. A., Walsh, J. K., Wamsley, E., Paulus, M., Kanady, J. C., & Drummond, S. P. A. (2013). The critical role of sleep spindles in hippocampal-dependent memory: A pharmacology study. *Journal of Neuroscience*, *33*(10), 4494–4504. https://doi.org/10.1523/JNEUROSCI.3127-12.2013

Minda, J. P., & Smith, J. D. (2002). Comparing prototype-based and exemplar-based accounts of category learning and attentional allocation. *Journal of Experimental Psychology:*

Learning, Memory, and Cognition, *28*(2), 275–292. https://doi.org/10.1037/0278-7393.28.2.275

Molle, M. (2006). Hippocampal sharp wave-ripples linked to slow oscillations in rat slow-wave sleep. *Journal of Neurophysiology*, *96*(1), 62–70. https://doi.org/10.1152/jn.00014.2006

Mölle, M., Bergmann, T. O., Marshall, L., & Born, J. (2011). Fast and slow spindles during the sleep slow oscillation: Disparate coalescence and engagement in memory processing. *Sleep*, *34*(10), 1411–1421. https://doi.org/10.5665/SLEEP.1290

Monaghan, P., Sio, U. N., Lau, S. W., Woo, H. K., Linkenauger, S. A., & Ormerod, T. C. (2015). Sleep promotes analogical transfer in problem solving. *Cognition*, *143*, 25–30. https://doi.org/10.1016/j.cognition.2015.06.005

Moscovitch, M., Cabeza, R., Winocur, G., & Nadel, L. (2016). Episodic memory and beyond: The hippocampus and neocortex in transformation. *Annual Review of Psychology*, *67*(1), 105–134. https://doi.org/10.1146/annurev-psych-113011-143733

Moscovitch, M., Rosenbaum, R. S., Gilboa, A., Addis, D. R., Westmacott, R., Grady, C., McAndrews, M. P., Levine, B., Black, S., & Winocur, G. (2005). Functional neuroanatomy of remote episodic, semantic and spatial memory: A unified account based on multiple trace theory. *Journal of Anatomy*, *207*(1), 35–66.

Murphy, G. L. (2016). Is there an exemplar theory of concepts? *Psychonomic Bulletin & Review*, *23*(4), 1035–1042. https://doi.org/10.3758/s13423-015-0834-3

Nadel, L., Samsonovich, A., Ryan, L., & Moscovitch, M. (2000). Multiple trace theory of human memory: Computational, neuroimaging, and neuropsychological results. *Hippocampus*, *10*(4), 352–368.

Ngo, H.-V. V., Martinetz, T., Born, J., & Mölle, M. (2013). Auditory closed-loop stimulation of the sleep slow oscillation enhances memory. *Neuron*, *78*(3), 545–553. https://doi.org/10.1016/j.neuron.2013.03.006

Nishida, M., Pearsall, J., Buckner, R. L., & Walker, M. P. (2009). REM sleep, prefrontal theta, and the consolidation of human emotional memory. *Cerebral Cortex*, *19*(5), 1158–1166. https://doi.org/10.1093/cercor/bhn155

Norman, K. A., Newman, E. L., & Perotte, A. J. (2005). Methods for reducing interference in the Complementary Learning Systems model: Oscillating inhibition and autonomous memory rehearsal. *Neural Networks*, *18*(9), 1212–1228. https://doi.org/10.1016/j.neunet.2005.08.010

Nosofsky, R. M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance*, *17*(1), 25.

O'Reilly, R. C., & Norman, K. A. (2002). Hippocampal and neocortical contributions to memory: Advances in the complementary learning systems framework. *Trends in Cognitive Sciences*, *6*(12), 505–510.

Oudiette, D., & Paller, K. A. (2013). Upgrading the sleeping brain with targeted memory reactivation. *Trends in Cognitive Sciences*, *17*(3), 142–149. https://doi.org/10.1016/j.tics.2013.01.006

Paller, K. A., Mayes, A. R., Antony, J. W., & Norman, K. A. (2020). Replay-based consolidation governs enduring memory storage. In M. S. Gazzaniga, G. R. Mangun, & D. Poeppel (Eds.), *The Cognitive Neurosciences 6th Edition* (6th ed.). MIT Press.

Pavlides, C., & Winson, J. (1989). Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *The Journal of Neuroscience*, *9*(8), 2907–2918. https://doi.org/10.1523/JNEUROSCI.09-08-02907.1989

Payne, J. D., Schacter, D. L., Propper, R. E., Huang, L.-W., Wamsley, E. J., Tucker, M. A., Walker, M. P., & Stickgold, R. (2009). The role of sleep in false memory formation. *Neurobiology of Learning and Memory*, *92*(3), 327–334. https://doi.org/10.1016/j.nlm.2009.03.007

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*(3, Pt.1), 353–363. https://doi.org/10.1037/h0025953

Rasch, J., Buchel, C., Gais, G., & Born, J. (2007). Odor cues during slow-wave sleep prompt declarative memory consolidation. *Science*, *315*(5817), 1426–1429. https://doi.org/10.1126/science.1134457

Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, *3*(3), 382–407.

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

Rudoy, J. D., Voss, J. L., Westerberg, C. E., & Paller, K. A. (2009). Strengthening individual memories by reactivating them during sleep. *Science*, *326*(5956), 1079–1079. https://doi.org/10.1126/science.1179013

Saletin, J. M., Goldstein-Piekarski, A. N., Greer, S. M., Stark, S., Stark, C. E., & Walker, M. P. (2016). Human hippocampal structure: A novel biomarker predicting mnemonic vulnerability to, and recovery from, sleep deprivation. *Journal of Neuroscience*, *36*(8), 2355–2363. https://doi.org/10.1523/JNEUROSCI.3466-15.2016

Schabus, M., Gruber, G., Parapatics, S., Sauter, C., Klösch, G., Anderer, P., Klimesch, W., Saletu, B., & Zeitlhofer, J. (2004). Sleep spindles and their significance for declarative memory consolidation. *Sleep*, *27*(8), 1479–1485.

Schapiro, A. C., McDevitt, E. A., Chen, L., Norman, K. A., Mednick, S. C., & Rogers, T. T. (2017). Sleep benefits memory for semantic category structure while preserving exemplar-specific information. *Scientific Reports*, *7*(1), 14869. https://doi.org/10.1038/s41598-017-12884-5

Schapiro, A. C., Turk-Browne, N. B., Botvinick, M. M., & Norman, K. A. (2017). Complementary learning systems within the hippocampus: A neural network modelling approach to reconciling episodic memory with statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *372*(1711), 20160049. https://doi.org/10.1098/rstb.2016.0049

Schechtman, E., Antony, J. W., Lampe, A., Wilson, B. J., Norman, K. A., & Paller, K. A. (2019). Multiple memories can be simultaneously reactivated during sleep as effectively as a single memory. *BioRxiv*. https://doi.org/10.1101/662015

Schönauer, M., Brodt, S., Pöhlchen, D., Breßmer, A., Danek, A. H., & Gais, S. (2018). Sleep does not promote solving classical insight problems and magic tricks. *Frontiers in Human Neuroscience*, *12*, 72. https://doi.org/10.3389/fnhum.2018.00072

Schouten, D. I., Pereira, S. I. R., Tops, M., & Louzada, F. M. (2017). State of the art on targeted memory reactivation: Sleep your way to enhanced cognition. *Sleep Medicine Reviews*, *32*, 123–131.

Schreiner, T., Göldi, M., & Rasch, B. (2015). Cueing vocabulary during sleep increases theta activity during later recognition testing. *Psychophysiology*, *52*(11), 1538–1543.

Schreiner, T., Lehmann, M., & Rasch, B. (2015). Auditory feedback blocks memory benefits of cueing during sleep. *Nature Communications*, *6*, 8729. https://doi.org/10.1038/ncomms9729

Schreiner, T., & Rasch, B. (2015). Boosting vocabulary learning by verbal cueing during sleep. *Cerebral Cortex*, *25*(11), 4169–4179. https://doi.org/10.1093/cercor/bhu139

Sio, U. N., Monaghan, P., & Ormerod, T. (2013). Sleep on it, but only if it is difficult: Effects of sleep on problem solving. *Memory & Cognition*, *41*(2), 159–166. https://doi.org/10.3758/s13421-012-0256-7

Smith, J. D., & Minda, J. P. (n.d.). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(6), 1411–1436.

Stark, S. M., Yassa, M. A., Lacy, J. W., & Stark, C. E. L. (2013). A task to assess behavioral pattern separation (BPS) in humans: Data from healthy aging and mild cognitive impairment. *Neuropsychologia*, *51*(12), 2442–2449. https://doi.org/10.1016/j.neuropsychologia.2012.12.014

Stickgold, R, Scott, L., Rittenhouse, C., & Hobson, J. A. (1999). Sleep-induced changes in associative memory. *Journal of Cognitive Neuroscience*, *11*(2), 182–193.

Stickgold, R, & Walker, M. (2005). Memory consolidation and reconsolidation: What is the role of sleep? *Trends in Neurosciences*, *28*(8), 408–415. https://doi.org/10.1016/j.tins.2005.06.004

Stickgold, R, & Walker, M. P. (2013). Sleep-dependent memory triage: Evolving generalization through selective processing. *Nature Neuroscience*, *16*(2), 139–145. https://doi.org/10.1038/nn.3303

Stickgold, Robert, & Walker, M. (2004). To sleep, perchance to gain creative insight? *Trends in Cognitive Sciences*, *8*(5), 191–192. https://doi.org/10.1016/j.tics.2004.03.003

Tamminen, J., Lambon Ralph, M. A., & Lewis, P. A. (2017). Targeted memory reactivation of newly learned words during sleep triggers REM-mediated integration of new memories and existing knowledge. *Neurobiology of Learning and Memory*, *137*, 77–82. https://doi.org/10.1016/j.nlm.2016.11.012

Tamminen, J., Payne, J. D., Stickgold, R., Wamsley, E. J., & Gaskell, M. G. (2010). Sleep spindle activity is associated with the integration of new memories and existing knowledge. *Journal of Neuroscience*, *30*(43), 14356–14360. https://doi.org/10.1523/JNEUROSCI.3028-10.2010

Tucker, M. A., Humiston, G. B., Summer, T., & Wamsley, E. (2020). Comparing the effects of sleep and rest on memory consolidation. *Nature and Science of Sleep*, *Volume 12*, 79–91. https://doi.org/10.2147/NSS.S223917

Tulving, E. (1972). Episodic and semantic memory. In *Organization of memory*.

van Dongen, E. V., Takashima, A., Barth, M., Zapp, J., Schad, L. R., Paller, K. A., & Fernandez, G. (2012). Memory stabilization with targeted reactivation during human slow-wave sleep. *Proceedings of the National Academy of Sciences*, *109*(26), 10575–10580. https://doi.org/10.1073/pnas.1201072109

van der Helm, E., Yao, J., Dutt, S., Rao, V., Saletin, J. M., & Walker, M. P. (2011). REM sleep depotentiates amygdala activity to previous emotional experiences. *Current Biology*, *21*(23), 2029–2032. https://doi.org/10.1016/j.cub.2011.10.052

Varela, F., Lachaux, J.-P., Rodriguez, E., & Martinerie, J. (2001). The brainweb: Phase synchronization and large-scale integration. *Nature Reviews Neuroscience*, *2*(4), 229–239.

Verleger, R., Rose, M., Wagner, U., Yordanova, J., & Kolev, V. (2013). Insights into sleep's role for insight: Studies with the number reduction task. *Advances in Cognitive Psychology*, *9*(4), 160–172. https://doi.org/10.5709/acp-0143-8

Wagner, U., Gais, S., Haider, H., Verleger, R., & Born, J. (2004). Sleep inspires insight. *Nature*, *427*(6972), 352–355. https://doi.org/10.1038/nature02223

Whitehurst, L. N., Cellini, N., McDevitt, E. A., Duggan, K. A., & Mednick, S. C. (2016). Autonomic activity during sleep predicts memory consolidation in humans. *Proceedings of the National Academy of Sciences*, *113*(26), 7272–7277. https://doi.org/10.1073/pnas.1518202113

Whitmore, N., Bassard, A. M., & Paller, K. A. (under review). Targeted memory reactivation of face-name learning depends on ample and undisturbed slow-wave sleep. *Npj Science of Learning*.

Wilson, M., & McNaughton, B. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, *265*(5172), 676–679. https://doi.org/10.1126/science.8036517

Witkowski, S., Schechtman, E., & Paller, K. A. (2020). Examining sleep's role in memory generalization and specificity through the lens of targeted memory reactivation. *Current Opinion in Behavioral Sciences*, *33*, 86–91. https://doi.org/10.1016/j.cobeha.2020.01.007

Yassa, M. A., & Stark, C. E. L. (2011). Pattern separation in the hippocampus. *Trends in Neurosciences*, *34*(10), 515–525. https://doi.org/10.1016/j.tins.2011.06.006

Zeithamova, D., Maddox, W. T., & Schnyer, D. M. (2008). Dissociable prototype learning systems: Evidence from brain imaging and behavior. *Journal of Neuroscience*, *28*(49), 13194–13201. https://doi.org/10.1523/JNEUROSCI.2915-08.2008

Zeithamova, Dagmar, Dominick, A. L., & Preston, A. R. (2012). Hippocampal and Ventral Medial Prefrontal Activation during Retrieval-Mediated Learning Supports Novel Inference. *Neuron*, *75*(1), 168–179. https://doi.org/10.1016/j.neuron.2012.05.010