

Mutual influences between memory and attention

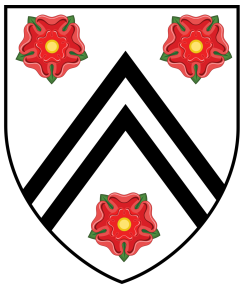
Nora Malika Roüast

New College

University of Oxford

DPhil in Experimental Psychology





Mutual influences between memory and attention

Nora Malika Roüast

Supervised by Professor Anna Christina Nobre & Dr Anna-Katharina Bauer

Submitted for the degree of DPhil in Experimental Psychology,

Hilary Term 2021

Abstract

Memory and attention are two of the most fundamental components of our cognition. Crucially, they do not operate independently but instead display an intricate, interactive relationship: our memories can guide our attention, and our attentional focus can shape our emerging memories. The aim of this doctoral thesis was to illuminate these mutual influences between memory and attention during the three crucial stages of memory establishment: encoding, consolidation, and retrieval. After delineating the scope of the thesis and defining relevant concepts and terminology in an introductory chapter (Chapter I), three experimental chapters discuss work exploring the memory attention interplay at these three stages.

In Chapter II, a distractor was used to divert attention during a learning task, thus targeting the encoding phase. The results showed a significant disruptive effect of the distractor during learning. Interestingly, distraction during encoding did not impair overall accuracy of subsequent memory retrieval and did not compromise the utility of the memories in guiding attention. However, more granular measures of memory

retrieval revealed systematic biases in the quality of memory representations. Electrophysiological markers also pointed to enduring effects of the distractor during encoding. These traces of the distractor at encoding suggest that inhibitory processes could have taken place but were not established enough to guide attention proactively or shape overall mnemonic accuracy.

In Chapter III, I attempted to modulate attention to a selective set of memories during consolidation by using targeted memory reactivations during a nap phase. During learning, searching for target locations in different categories of scenes was linked to distinctive tunes, which were novel and purposely composed for the study. Participants then had a nap opportunity during which electroencephalography was recorded. A subset of tunes was replayed during napping periods when the participant was judged to be in slow-wave sleep to reactivate spatial contextual memories linked to those particular scene categories, while tunes associated to the other categories were not replayed. Later in the day participants performed memory retrieval and memory-based orienting tasks to compare performance on replayed vs. non-replayed scene categories. The findings diverged from the frequent observations of sleep-based benefits on memory strength in the sleep literature. Indeed, sleep-based stimulation could even lead to potential detriments of the memory as attentional guide. The implications of these findings, and their reliability, are discussed.

The set of studies described in Chapter IV manipulated selective attention to contents within memories at the stage of memory retrieval. The approach of using cues to direct attention retroactively to memory content developed within the working-memory literature (retrocues) was adapted for use with long-term memory retrieval. Spatially informative retrocues were found to speed and improve the identification of memory targets as well as, given the right conditions, the spatial precision of the

memory. To my knowledge, this comprises the first evidence for the efficacy of retrocues to investigate selective retrieval in long-term memory. These observations imply that the selective retrieval of long-term memories is susceptible to attentional modulation. The newly developed experimental approach shows promise for adaptation with advanced research methodologies, for different populations, and for various lines of investigation.

Acknowledgements

The story behind a doctorate is never about just one person but is populated by a host of people providing support, wisdom, skills, and fun. I was incredibly fortunate to have a wonderful team by my side that stuck with me, provided me with countless opportunities, and helped me through this turbulent time.

First and foremost, I am simply so grateful to my wonderful supervisor Kia Nobre. Aside from her brilliant scientific mind and generous, fun nature, she has shown true empathy and kindness to me for all these years. She really has gone above and beyond to support me through whatever life chose to throw at me. I cannot put into words what her enduring kindness and support mean to me and this doctoral thesis would not exist without her. There is definitely a reason why Germans use the description of "doctoral mother" for supervisors.

I was indeed very lucky to have a team of strong female mentors, teaching me a wide variety of skills and sharing their wisdom. I sure tried to soak up their valuable lessons on life and science as quickly as I could, yet I know that it must have taken a lot of patience and good will to provide them again and again. So, in chronological order, I thank all of you: Zita, Ana, Katharina, and Nahid. I was honoured to work with you, and I will treasure the memories of my time with you – be it learning new methods, fixing codes, fun at conferences, or just some direly needed chat and catch up. Katharina deserves a special highlight for taking over as my co-supervisor and literally picking up the pieces of my research, providing sustenance through long testing sessions, and maintaining a humorous take on every disaster on the way.

I would also like to thank Prof Mark Stokes, Prof Paul Dux, and Prof Emily Holmes. Each provided me with opportunities and support at different stages of my thesis. Mark acted as my initial co-supervisor and provided valuable feedback. Paul let me visit his lab at the University of Queensland in Brisbane, Australia. And Emily invited me for an overseas institutional visit to the Karolinska Institute in Stockholm, Sweden, where I was able to learn much about trauma and PTSD. All of these, indeed any of my research would have been impossible without my funding body: I was extremely lucky to be funded by the ESRC and I have nothing but good memories of and gratitude for their role in my DPhil studies.

In general, I was very fortunate to stumble into an incredible lab: throughout different generations and times, the Brain & Cognition lab was always an amazing group of people. They were and are a constant source of inspiration, fun, kindness, and support, so much so that they became my friends rather than colleagues. All of you are wonderful and I wish you the best. Specific shoutouts go to Giedrė, Lauren, and Simone for the great girl chats, Sammi for his unfailing ability and willingness to solve every one of my problems (you are a star!), Irene and Katharina for sleep scoring, Jasper for his care in difficult times and swimming sessions, Sarah and Yvonne for post-baby magic help, and of course Rob Mok Soc.

On a more private note, a thank you to Hertford Choir and Katie for keeping a song on my lips and in my heart, and to Jellie for her unwavering belief in me and her friendship that kept me sane. To Lizzie for her creativity and amazing skills as a soundboard. To Lewis for his music. To Lisa for fixing my ailments and discussing life and the world with me. And to Oxford's beauty for making me look up in wonder again and again.

Especially the last year I would not have managed without the help of friends and family. Thanks to my granny, without whom none of this would have even begun. To my Dad and Patty, and Birgit and Andreas for opening their homes to us. To my Mum for her untiring support of us. To my brothers Philipp and Onno for chats and visits around the world. And to Eva and Timo for always, just always, having our backs when we needed it and standing by our sides (even as marriage witnesses) with a solution for everything. I will never forget. To Hekkie for his cuteness and chonky cuddles. And, of course, to Alan, the love of my life, for being the best person I could imagine. I am so glad you asked me about fire poi many years ago.

Contents

Abstract **2**

Acknowledgements **5**

Contents **8**

Chapter I. Introduction **11**

Memory **11**

Attention **15**

Memory and Attention in concert **19**

Memory guidance of attention **20**

Attentional influences on memory **22**

… Encoding 22

… Consolidation 23

… Retrieval 25

Clinical relevance of studying memory and attention interactions **27**

Influencing memory and attention interactions **28**

At encoding: Distraction 30

At consolidation: Targeted memory retrieval 30

At retrieval: Retrocues 31

In the light of the literature 31

Chapter II. Distraction during encoding **33**

Introduction **33**

Method **43**

Participants 43

Materials 44

Design 46
Procedure 52
Behavioural Analysis 54
EEG Data Acquisition 56
EEG Analysis 57
Results 62
Learning target locations with and without distraction 62
Recalling target locations with and without prior distractor association 63
Lingering distraction effects on established long-term memory 66
Memory-based attention with and without prior distractor association 67
Distraction association modulates the neural response to memory cues 69
Distraction association modulates effect of memory on target processing 70
Discussion 74
Chapter III. Selective reactivation during consolidation 83
Introduction 83
Method 91
Participants 91
Materials 92
Design 96
Procedure 101
Behavioural Analysis 104
EEG Data Acquisition 107
EEG Analysis/Sleep scoring 108
Results 113
Behavioural results 113
Morningness-Eveningness Questionnaire results 122
Sleep-dependent results 122

Stimulation-dependent results 126

Discussion 129

Chapter IV. Selective attention during retrieval 140

Introduction 140

Study 1: Developing an experimental approach to investigate selective attention during LTM retrieval 146

Method 146

Results 155

Discussion 160

Study 2: Replicating benefits of selective attention in LTM retrieval 166

Method 166

Results 169

Discussion 175

Study 3: Comparing consequences of attending and ignoring items in LTM . 177

Method 178

Results 180

Discussion 185

Overall discussion 186

Chapter V. General discussion 201

Distraction during memory encoding 201

Targeted memory reactivation during memory consolidation 208

Retrocues during memory retrieval 220

The way forward 227

Concluding remarks 228

References 230

Chapter I. Introduction

Over three experimental chapters, I will explore mutual influences between memory and attention. Memory and attention are two of our main cognitive functions (Wixted, 2018), having inspired decades or even centuries of scientific investigation. I will set out by discussing relevant concepts and terminology for each of them, thus defining the scope of this thesis. Subsequently, I will demonstrate that these two domains work interdependently, with memory guiding attention and attention modulating memories. Importantly, these mutual influences appear to occur throughout the three essential stages of memory: encoding, consolidation, and retrieval. My experimental work explores the interplay between memory and attention, after a modulation at each of these levels: distraction during encoding (Chapter II), targeting selected memories for preferential consolidation (Chapter III), and biasing retrieval using retrocues (Chapter IV).

Memory

"*Mental states of every kind, -- sensations, feelings, ideas, -- which were at one time present in consciousness and then have disappeared from it, have not with their disappearance absolutely ceased to exist."* (Ebbinghaus, 1885)

The study of the mental states that Ebbinghaus describes, the memories we accumulate over our lifetime, has long captivated scholars. Consequently, a breadth of conceptual distinctions, terminology, and bodies of research have come to be considered, in giving structure and providing a conceptual framework to the study of memory. The scope of what the term "memory" encompasses is undeniably large:

memory spans from the lingering of impressions on our senses (iconic memory), to keeping relevant information in mind while completing a task (short-term or working memory, WM), to childhood reminiscences spanning many decades (long-term memory, LTM, Atkinson & Shiffrin, 1968). The following studies in this thesis almost exclusively focus on long-term memory, specifically contextual and associative episodic-like long-term memory: Episodic memory (Tulving, 1983, 2002) allows us to retrieve personal experiences explicitly, often consisting of a conglomerate of identity, spatial, and temporal information (the "what", "where", and "when", Clayton & Dickinson, 1998).

Neurological observations of memory loss in the case of amnesia (famously the case of H.M. comes to mind, Scoville & Milner, 1957; Corkin, 2013) and dementia (namely Alzheimer's Disease, Kolb & Wishaw, 1996) have emphasised the importance of the hippocampus and associated medial-temporal structures for functional memory processing. Behavioural and neural studies have confirmed a crucial role of the hippocampus for episodic memory (Moscovitch et al, 2016), often with a specific emphasis on spatial memory (Burgess et al, 2002).

These neurological studies are not only informative about the underlying neural substrate that supports memory, but also for conceptual distinctions within memory. The case of H.M., for example, allows a separation of semantic ("facts") and episodic memory, placing emphasis of hippocampal function on the latter category. The hippocampal resection in H.M., a treatment for intractable epilepsy, led to severe, persisting anterograde amnesia, leaving him unable to form new episodic memories (Scoville & Milner, 1957; Corkin, 2002). However, whilst he showed a degree of retrograde amnesia and struggled to recall episodic memories from the two years leading up to the operation, his memory for remote episodic information was relatively

spared (Scoville & Milner, 1957; Milner et al, 1968). Observations like these informed the conceptual separation of memory processing into encoding, consolidation, and retrieval: H.M. was unable to encode new, yet able to retrieve older episodic information. The presence of retrograde amnesia suggests there must also be an intermediate stage, during which originally hippocampal-based memories are consolidated and become potentially independent of an intact hippocampus. These three components, encoding, consolidation, and retrieval, are commonly referred to as the three main processing stages of memory (for example, Wilhelm et al, 2011a; Westermann et al, 2015; Crowley et al, 2019).

Encoding is a necessary precursor for successful remembering: It denotes the process by which we register, learn or combine new information, often based on perceptual input. For episodic memory only a single exposure can be sufficient to form complex memories spanning temporal to spatial, visual and auditory, as well as contextual information. However, as with any learning mechanism, a balance has to be struck between maintaining learnt information and updating by integrating novel information ("stability-plasticity dilemma"): Stored memories must be maintained without blocking new encoding, but at the same time without being overwritten by the novel input ("catastrophic interference"). According to the two-stage model of memory (Marr, 1971; McClelland et al, 1995), the human memory system achieves this balance with a fast, yet temporary storage system and a secondary slow, yet stable store.

The progressive stabilisation of long-term memory by transferring the memories from the fast to the stable storage system is termed consolidation (Dudai, 2004). Even though the nature of the consolidation process and the resulting memories is debated (Winocur et al, 2010), there is general agreement about the occurrence of consolidation and the beneficial impact of sleep (Maquet, 2001; Stickgold, 2005; Born et al, 2006;

Diekelmann & Born, 2010). For the sake of completeness, the emergence of evidence for memory consolidation during wakeful rest should be mentioned (see Wamsley, 2019, for discussion). Within this thesis I shall focus on sleep-based consolidation. The Active Systems Consolidation hypothesis stipulates that this sleep-based benefit on consolidation is not merely a result of time spent without interference (Ellenbogen et al, 2006) but that sleep physiology actively promotes memory consolidation (Born & Wilhelm, 2012; Inostroza & Born, 2013). Notably, there might be multiple ways in which sleep-based consolidation could take place (Schönauer et al, 2015). In general, the transfer and thus consolidation of memory is thought to occur via hippocampal-based reactivation of memories (Wilson & McNaughton, 1994). Sleep-based reactivation of memories does not only occur spontaneously but can also be triggered by olfactory (Rasch et al, 2007) or auditory (Rudoy et al, 2009; Bendor & Wilson, 2012) cues: the technique of targeted memory reactivation uses the principle of pattern completion (e.g., Rolls, 2013) to trigger the reinstatement of memories by presenting the sleeping participant with an associated cue or part of the memory. Cueing sleep-based reactivation and presumably associated consolidation have been widely shown to benefit memory (for a meta-analysis of studies using targeted memory reactivation, see Hu et al, 2020).

The reactivation discussed consists of patterns of neural activity during wakefulness replaying during sleep, for example in the same temporal sequence (Skaggs & McNaughton, 1996), yet compressed in time (Nádasdy et al, 1999). This reactivation bears similarities to the process of retrieving a memory: Retrieval refers to the recollection of the memory, which reinstates the patterns of brain activity experienced during encoding (Danker & Anderson, 2010). This observation gives new credence to the idea of memory retrieval as a form of mental time travel (Tulving,

2002). Similarly to the reactivation during sleep described above, memory retrieval can also be triggered by external or internal cues, giving rise to the reactivation of the associated hippocampal and neocortical networks (Moscovitch et al, 2016). This retrieval process is thought to be separated into two stages (Moscovitch, 2008): Firstly, the cue interacts rapidly and unconsciously with the hippocampus ("ecphory"), and consequently the neocortical traces. Secondly, conscious and slower cortical processes act on the resulting output bringing about the conscious experience. Notably, Waldhauser and colleagues (2016) show that ecphory functionally drives the episodic memory retrieval. The idea that retrieved memories are indeed reinstated is supported by evidence showing effects of retrieval on the memory, such as memory benefits (Carrier & Pashler, 1992) and forgetting of competing memories (Wimber et al, 2015). Amongst others, these findings suggest that memories are malleable and that their influence is not limited to the memory itself.

Attention

The second cognitive function this thesis focuses on is attention. Despite attention permeating through almost all psychological phenomena, it is famously difficult to describe or define. Time and time again, researchers have helped themselves to the often-quoted William James (1890) to grapple with the concept of attention as "*the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence.*" In this description, James touched on multiple important aspects of attention, including the idea of selection. Selective attention refers to the process of singling out specific events, items or thoughts while ignoring irrelevant competitors in the environment. By this, selective attention can be defined as the "*set*

of functions that prioritize and select what comes to prevail in mind" (Nobre, 2018). The process by which we are able to achieve this selection has been the subject of many classic theories (e.g., Broadbent, 1958; Treisman, 1969). Dominant paradigms for the study of selective attention have been visual search tasks (see Müller & Krummenacher, 2006), tasks with distraction elements (see Chelazzi et al, 2019), or tasks using attentional cues (see Posner, 1980). Within her definition, Nobre (2018) also emphasises a different key aspect of selective attention as being a "*proactive process, which prepares and guides perception in an anticipatory way*". Evidence for the proactive nature of selective attention can be seen, for example, in the facilitation of processing afforded to locations of likely relevance (Kastner et al, 1999; Stokes et al, 2009). Attention within the context of this thesis will solely focus on the area of selective attention, while acknowledging the existence and importance of other forms of attention such as sustained attention ("concentration").

A further important distinction within the field of attention is the framework of externally and internally oriented attention (Chun et al, 2011). Attention allows us to modulate and select information based on sensory input ("external") and internally generated information such as long-term memory ("internal"). Wundt (1874) came close to this distinction by describing the two main psychological processes bound to consciousness as "*der eine ist die Bildung von Vorstellung aus Sinneseindrücken, der andere das Gehen und Kommen der reproducierten Vorstellungen*" (sic, chapter 18, p. 711). This roughly translates to "the one is the production of ideas from sensory impressions, and the other the goings and comings of the reproduced ideas", describing awareness for ideas of both external and internal origins. This distinction into external and internal can apply to what drives the attentional focus as well as to the substance on which attentional selection acts (Nobre, 2018).

Firstly, the influences driving attention can be both of external and internal origin (Wundt, 1874: ”*Die Einflüsse, welche die Apperception lenken, sind theils äussere theils innere*.”, sic, chapter 18, p.720). From an evolutionary perspective, we need to be able to react, often rapidly, to external information such as danger or the presence of food while foraging. Physical salience in the environment is thought to drive the capture of “bottom-up” attention, which occurs fast and arguably involuntarily, as a “physiological reflex” (Posner, 1980). On the other hand, humans also had to develop a system by which we could complete tasks, requiring orienting attention according to an internally generated set of task goals. “Top-down” attention (Theeuwes, 2018) is supposed to allow volitional orienting of attention to achieve our aims, without unduly being distracted by irrelevant information or events. The idea that our attention is solely guided by a combination of involuntary, external and voluntary, internal biases (Jonides, 1981) is now thought to be somewhat incomplete. More recently, Awh and colleagues (2012), amongst others, proposed that attention could also be guided by selection history, comprised of internal factors. Nobre and Stokes (2019) argue to extend this definition to include a variety of mnemonic information as potential sources of attentional guidance. Experimentally, external and internal influences or biases on attention have been manipulated by the use of cues, the efficacy of which has been replicated in multiple studies. Famously demonstrated through the Posner cueing paradigm (Posner et al, 1978; Posner, 1980), external cues in form of arrows can bias our perception and hence our spatial selective attention in the direction of the arrow. Internal, voluntary biases have also been demonstrated to bias attention, even overriding bottom-up capture (Leber & Egeth, 2006). Furthermore, memory-guided attention tasks have proven that we also use internally generated information, such as

long-term memories, to guide our spatial attention (Summerfield et al, 2006; Stokes et al, 2012; Rosen et al, 2014).

Secondly, the substances on which attentional guidance acts are proposed to be the external, sensory world, as well as the internal, mental world (Chun et al, 2011). Indeed, significant similarities have been reported in the way we use attention to orient in the external world as well as the internal world: We can use our spatial attention to search an apple tree to select the ripe fruits, mirrored experimentally by visual search tasks (Müller & Krummenacher, 2006), and we can also use our spatial attention to search our memory (Nordfang & Wolfe, 2018). This has already been observed by Wundt (1874), who wrote on how our attention or "inner focus-point" can focus successively on different parts of this internal field of vision (verbatim translation of: "*Der innere Blickpunkt kann sich nun successive den verschiedenen Theilen des inneren Blickfeldes zuwenden.";* sic, chapter, 18, p.718). Further, both internal and external worlds contain too much information to be processed all at once, introducing the need for attentional selection (Chun et al, 2011), as described above.

Attentional selection can be demonstrated both amongst perceptual input (Treisman, 1969) and amongst memories (Wais et al, 2010; Wimber et al, 2015; Tarder-Stoll et al, 2020). The introduction of retrocues was an important motivation for the study of selective attention within mnemonic space (Griffin & Nobre, 2003; Landman et al, 2003): Griffin and Nobre observed that retrospectively informative cues (hence, *retrocues*) significantly benefitted participants' working memory performance: Participants were faster and more accurate in identifying whether an object had been part of an array if they were retrospectively informed about the spatial location of the to-be-probed item. Within various working memory paradigms, the efficacy of retrocues has been replicated widely (review Souza & Oberauer, 2016; Nobre & Stokes,

2020), and there is evidence that retrocues might influence both access to and precision of the cued memorandum (Makovski & Pertzov, 2015; Gunseli et al, 2015; Myers et al, 2018). Whilst the functional mechanisms through which retrocues act are still debated (Souza & Oberauer, 2016; Myers et al, 2017), the similarities between attentional selection in external and internal, mnemonic space have been noted (Gazzaley & Nobre, 2012). Indeed, within the context of long-term-memory research, the influential Attention To Memory (AToM) theory stipulates that the internal orienting of attention during retrieval might mirror the attentional selection observed for the external world (Ciaramelli et al, 2008). Specifically, AToM proposes a crucial role of the parietal cortex in internal orienting, with functionality divided along goal-based attention and attentional capture in parallel to the proposal for perceptual attention by Corbetta and Shulman (2002).

Memory and Attention in concert

The observation that attentional selection operates within memory, and indeed the existence of theories like AToM, highlight that memory and attention are not separate cognitive entities, but might interact intimately to guide performance. Chun and Turk-Browne (2007) argue for an interdependence of the two cognitive functions based on conceptual issues and neural observations: Firstly, both attention and memory present with limited capacity, arguably addressed by attention deciding what information gets encoded (Turk-Browne et al, 2006) and memory from past experiences guiding attention (Summerfield et al, 2006). Dudukovic and Wagner (2006) eloquently summarise this interaction by highlighting how "*attention and memory work in concert, with attention influencing which memories are formed and memory guiding how visuospatial attention is allocated*". Secondly, this interaction is further supported by

the observation that attentional networks are recruited in memory tasks and memory areas are influenced by attentional states: Summerfield and colleagues (2006) found that memory-guided attention involved the parietal-frontal network, known to be important for visual spatial orienting. Aly and Turk-Browne (2016a, 2016b) in turn observed that attention stabilised hippocampal memory representations, thereby aiding memory consolidation. Consequently, these interactions between memory and attention challenge typical perceptions of each function, showing that memory is not only retrospective and attention is not solely proactive (Nobre & Stokes, 2019, 2020).

Memory guidance of attention

A growing field of scientific interest is focused on how memories can proactively influence our perception by preparing and guiding attention. From an evolutionary standpoint, it is unsurprising that information about our environment collated in episodic memories would be used to our benefit, for example by guiding our attention in visual search (Hollingworth, 2006). Even if the memories are not explicitly retrievable, previous exposure to spatial layouts helps identifying locations of interest (contextual cueing, Chun & Jiang, 1998; Chun, 2000). However, memories are not only limited to providing information about future actions or decisions, they have been hypothesised to generate predictions about the anticipated future (Bar, 2009) and thereby to be proactively involved in preparing our perception and behaviour (Nobre & Mesulam, 2014; Nobre & Stokes, 2019, 2020). Nobre and Stokes (2019) discuss how this "*premembering*" can occur over different timescales: in the very short term, the impressions of prior visual input capturing attention can enhance following stimuli by creating an excitatory state (Posner, 1980). At a somewhat larger time scale, there is general agreement of the role of working memory in guiding attention (Desimone and

Duncan, 1995), be this voluntary guidance according to task aims or involuntary (Carlisle & Woodman, 2011; Soto et al, 2005), or as a transfer resulting from previous input or choices (e.g., Maljkovic & Nakayama, 1994). Studies of memory-guided attention (Summerfield et al, 2006; Stokes et al, 2012) and change-blindness (Rosen et al, 2014; Rosen et al, 2015a, 2015b) confirm that even long-term memories can guide our, in these studies visuo-spatial, attention to locations of likely importance.

Apart from studies showing these behavioural impacts, as well as perceptual sensitivity (Patai et al, 2012) and guidance of eye movements as marker of attention (Hannula et al, 2010; Meister & Buffalo, 2016), there are also various studies that document neural markers of attentional guidance by memory. Memory cues led to a modulation of activity levels in the hippocampus (Summerfield et al, 2006; Goldfarb et al, 2016; Günseli & Aly, 2020; but see Rosen et al, 2018), yet also engaged mechanisms of preparatory spatial attention such as spatially specific desynchronization of alpha oscillations (e.g., Stokes et al, 2012). In response to memory-guided attention, further studies found modulation of event-related potentials (ERP, Summerfield et al, 2011; Doallo et al, 2013), an impact on attentional visuospatial representations (Rosen et al, 2015a), as well as the activation of attention and control networks (Rosen et al, 2015b). Indeed, a recent meta-analysis on memory-guided attention (Fischer et al, 2021) stressed the involvement of the dorsal fronto-parietal attention network and the fronto-parietal control network for memory-guidance of attention. After reviewing rodent and human data, Leszczyński and Staudigl (2016) proposed that information from and communication between hippocampal memory and fronto-parietal attention might be achieved via the anterior thalamus, as a connecting and modulating hub allowing for memory-guided attention. The proposal of a functional role of the thalamus for memory-guided attention is particularly interesting given the importance of the

thalamus, and thalamic spindles, for the consolidation of memory (Rasch & Born, 2013).

Attentional influences on memory

In addition to the observation that memories have sway over the allocation of attention, it is known that the focus of our attention influences the formation (and content) of the memory itself. Kuhl and Chun (2014) argue that the interplay of memory and attentions goes beyond individual instances, instead preferring to frame memory as acts of selective attention: "*The ways in which we form, retrieve, and work with our memories largely represent acts of attention*". Beyond the implication of memory as being attentional acts, this idea implies that attention could influence memory at the different stages of establishment discussed above. Indeed, effects of attention have been demonstrated for encoding, consolidation, and retrieval of memories.

… Encoding

Attention is often considered a prerequisite for successful encoding, and later remembering, with classic studies such as the dichotic listening task (Cherry, 1953) showing virtually no memory of the unattended auditory input (Moray, 1959). Further, dividing attention between two tasks, thus allocating less attention to the information in question, impairs the formation of new memories (Fernandes & Moscovitch, 2000; for review see Aly & Turk-Browne, 2017). Even the presence of a distractor at encoding, such as a social stimulus, and the resulting diversion of attention impacts subsequent memory acuity negatively (Doherty et al, 2017). Rather than relying merely on the memory task by itself, subsequent-memory paradigms also allow observing the impact of selective attention during and on memory encoding, informed by whether a

memory was ultimately formed (for discussion, see Kuhl & Chun, 2014). Different attentional modulations can be observed during encoding associated with memory success or failure depending on distinct activation patterns or the directionality of attention (Uncapher & Wagner, 2009; Uncapher et al, 2011). Importantly, Uncapher and Rugg (2009) observed the attentional modulation of hippocampal activity to be selective for the attended feature of the objects. In general, research has shown that attended stimuli are associated with stronger neural representation and better subsequent memory accuracy, when compared to competing, unattended stimuli (Johnson & Zatorre, 2005; Yi & Chun, 2005; see Kuhl and Chun, 2014). Given the emphasis of the thesis on hippocampal memories, it is important to note that the attentional state of the hippocampus during encoding served as a predictor for memory accuracy, showing stronger representations for subsequently remembered information (Aly & Turk-Browne, 2016a).

… Consolidation

Assessing attentional effects on the consolidation of memories, especially during sleep, is somewhat less straightforward given the ordinarily antithetical nature of attention and sleep. It is possible, however, to approximate attentional effects on consolidation indirectly. Consolidation has been shown to be selective to memories that are expected to be of future relevance (Wilhelm et al, 2011b), thus putatively occupying a higher attentional status. Further, memories associated with anticipated reward show higher learning during sleep (Fischer & Born, 2009). Given that reward is associated with preferential attentional selection (Anderson, 2016; Tibboel & Liefooghe, 2020), it is conceivable that the reward-associated memories were consolidated because of this "attentional tag".

Further indication that attention could influence the consolidation process stem from rodent animal models, specifically experiments observing the thalamic circuitry: the thalamus appears to have a crucial role in memory consolidation during sleep and in regulating attentional selection during wakefulness (for discussion, see Klinzing et al, 2019). The neurons in question are specifically GABAergic neurons of the thalamic reticular nucleus (TRN) which in interplay with thalamocortical neurons generate spindle oscillations (e.g., Steriade et al, 1993). Thalamic spindles in humans are thought to be the driving force of sleep-dependent memory consolidation (Rasch & Born, 2013; Latchoumane et al, 2017; Antony et al, 2019). In addition, TRN neurons have been found to modulate sensory selection in wakefulness if attention is divided (Wimmer et al, 2015). Interestingly, the same TRN neuronal populations were associated with both sleep-based spindles and attention-related alpha oscillations in mice (Chen et al, 2016). The TRN neurons also appear to be divided into subnetworks depending on sensory or limbic projection, showing dissociative patterns of synchrony and suppression with sleep physiology and attention (Halassa et al, 2019). Apart from linking consolidation and attentional processes, Klinzing and colleagues (2019) speculate that this common circuitry might permit information to receive an attentional tag for subsequent sleep-based consolidation.

The experimental paradigm of targeted memory reactivation (Paller et al, 2021, also see above) gives an interesting perspective on the relationship between attention and memory consolidation. Targeted memory reactivation uses a fundamental element of a memory (a smell, Rasch et al, 2007; or sound, Rudoy et al, 2009) as a cue to trigger memory reinstatement during sleep (Bendor & Wilson, 2012; Staresina & Wimber, 2019). Just as cues guiding our external and internal attention (Chun et al, 2011), could this memory cue be seen as guiding selective retrieval, and thus an act of attention

(Kuhl & Chun, 2014)? Of course, one needs to acknowledge that participants are asleep and by definition unaware of the memory cue, showing no signs of arousal from sleep. Nonetheless, accepting this drawback, it is intriguing to think of the memory cues as involuntarily and unconsciously guiding internal attention in the sense of prioritising and selecting memories for consolidation. The drawback of lacking awareness during sleep might not completely preclude evaluating the relationship between attention and consolidation: Recently, the scientific study of lucid dreamers has received interest (Konkoly et al, 2020; Paller et al, 2021). Lucid dreamers are individuals who can be aware of themselves and that they are sleeping during the REM sleep phase (Baird et al, 2019). Konkoly and colleagues have been successfully attempting to communicate with lucid dreamers by eliciting previously discussed answers in the shape of eye movements from the lucid dreamers by using verbal cues from the experimenter. In principle, the existence of lucid dreamers, and communication with them, substantiates the possibility of attentional effects on consolidation and opens new avenues of studying this putative relationship.

… Retrieval

According to a proposal by Kuhl and Chun (2014) attention does not merely influence memory retrieval; instead, the retrieval of the memory itself can be considered an act of selective attention. These considerations are supported by observations that selective retrieval involves similar neural networks as orienting our selective perceptual attention (Ciaramelli et al, 2008; Cabeza, 2008; Cabeza et al, 2008). Regardless of the conceptual framing, the importance of attention for memory retrieval cannot be understated (Aly & Turk-Browne, 2017), with Tarder-Stoll and

colleagues (2020) identifying "*attentional selection of goal-relevant memories*" as one of three key components of Tulving's hypothesised retrieval mode (Tulving, 1983).

The effects of divided attention were found to be less pronounced for retrieval than for the initial encoding of memories (Craik et al, 1996). Decrements in performance are often restricted to tasks that divide attention within the same domain, such as the verbal domain (Fernandes & Moscovitch, 2000). However, the effects of selective attention on retrieval are indubitable, showing attentional modulation of supporting hippocampal activity during retrieval (Hashimoto et al, 2012; Duncan et al, 2012) and strengthening of selected memories, while weakening competing ones (retrieval-induced forgetting; Wimber et al, 2015). The effects of competing memories, and by association the impact of successful attentional selection, could also be observed in reactivation in the ventral temporal cortex: Kuhl and colleagues (2011, 2012) found that competing memories were not only reactivated, but the strength of this reactivation was associated with worse memory performance. Further, the presence of a distractor during retrieval, detracting from the attentional focus and taking up underlying neural resources, was also associated with reduced memory-related signals in the hippocampus and impaired memory performance (Wais et al, 2010).

More specifically, the study of retrocues gives an insight into the role of selective attention for memory retrieval. As discussed above (section Attention) retrospectively informative cues (retrocues) point observers to a specific aspect of their memory, be it a spatial location, a feature, or an item, thus guiding internal attention to the relevant mnemonic content (for review, see Nobre & Stokes, 2020). Consequently, the selected memory content highlighted by the retrocue and thereby attention, has been associated with greater recall speed (Griffin & Nobre, 2003) and precision (Makovski & Pertzov,

2015; Gunseli et al, 2015), whilst the unattended or selected-against memory items degraded in quality (LaRocque et al, 2015).

In general, the efficacy of retrocues speaks to the active role of internally guided attention on the retrieval of memories. Until now these effects have only been observed in working memory, with some researchers arguing that retrocues selectively benefit the maintenance of the retrocued items (Pertzov et al, 2013). The effects on pure retrieval of long-term memory have not been previously shown. One of the aims of this thesis is to explore the applicability of retrocue manipulations to investigate selective attention during long-term memory retrieval (see Chapter IV. Selective attention during retrieval).

Clinical relevance of studying memory and attention interactions

As the previous section details, it is apparent that the functionality and underlying neural systems of memory and attention are not as separate as perhaps once thought. The connectivity of memory and attention, the mutual influences on each other, have many real-world applications. To name a few, attentional biases towards negative, emotional, or threatening events or memories occur in depression and anxiety (Dalgleish & Watts, 1990; Gotlib et al, 2004; Gotlib & Joorman, 2010), and in post-traumatic stress disorder (Uddo et al, 1993; Neylan et al, 2004; Bomyea et al, 2017). Indeed, multiple behavioural treatment approaches focus on the relationship of memory and attention: Cognitive bias modification for emotional disorders such as depression is aimed to change attentional bias from negative to positive, be it external input or internal information and memories (e.g., Koster et al, 2009). The gold standard for treating post-traumatic stress disorder, cognitive behaviour therapy (e.g., Foa & Rothbaum, 2001; National Collaborating Centre for Mental Health (UK), 2005), works

with the traumatic memories in order to reduce their negative impact on cognition, memory and attention. Interestingly, and importantly, in a review on cognitive bias modifications in post-traumatic stress disorder, Woud and colleagues (2017) report that different treatment protocols target different biases, with focus on memory or attention, showing varying levels of successful outcomes. Plainly, investigating how influences act on and modulate memory and attention, and at what stages the interaction is susceptible to input, has direct clinical impact.

Further, there are questions about the stage of the memory and its malleability that have direct impact on treatment options: Early treatment attempts of post-traumatic stress disorder specifically target the intrusive, unwanted retrieval of traumatic memories ("flashbacks", a hallmark of post-traumatic stress disorder (DSM- IV-TR; American Psychiatric Association, 2000): The "Tetris" approach, targeting the consolidation of traumatic memory, aims to reactivate traumatic memories, leaving them labile, and then to engage visuospatial memory resources with a competing task – Tetris (Holmes et al, 2009). This approach is aimed at, and appears successful in, alleviating or even preventing intrusive flashbacks from occurring, which otherwise force memory retrieval and "reliving" of traumatic memories and "hijack" attentional resources (James et al, 2015). It has been shown to hold promise in reducing traumatic memories after real-life traumatic events associated with motor-vehicle accidents (Iyadurai et al, 2018).

Influencing memory and attention interactions

Given the wider implications and clinical importance of interactions between memory and attention, it is of utmost importance to understand how memory and attention interact throughout the establishment of memory. At what points is the

memory malleable for different attentional impacts? How do differing influences at encoding, consolidation, and retrieval impact on the resulting memory and the use of the memory as an attentional guide? In general, what are mutual influences that shape our memory and attention?

Considering the complexity of studying not only two different cognitive domains, but also their interactions, it is important to find a unifying approach to address these arising research questions. Whilst previously discussed studies often show strength in their simplicity of focusing solely on, for example, consolidation of selected memories by future relevance (Wilhelm et al, 2011b), or distraction effects on memory retrieval (Wais et al, 2010), assessing mutual influences of memory and attention requires a multi-tiered paradigm. Memories have to be formed and encoded in the first place, before one can assess their potential for and qualities in both attentional guidance and retrieval acuity. All experimental paradigms in this thesis thus include three main components: Firstly, all experiments include a learning task, in which episodic memories can be acquired in a controlled environment. Specifically, I use visual search in complex naturalistic environments (Biederman, 1972), in which targets (or distractors) have been embedded. Over subsequent learning blocks, location-scene associations are formed, potentially supplemented by further contextual information such as sound or distractors. Secondly, every experiment requires at least one memory task, implemented as a direct recall task, in which participants have to identify the identity and location of the target within the scene context. Thirdly, each study includes an attention task, in which the attentional selection guided by or occurring within memories can be evaluated. Depending on the conceptual motivation for the task, I either used modifications of the memory-guided attention task introduced by

Summerfield and colleagues (2006), or adapted the retrocueing manipulation (Griffin & Nobre, 2003) for use with long-term memories.

The multiple phases of the experimental studies allowed manipulating the relationship of attention and memory at each stage of memory establishment and assessing the resulting influences on memory and attention. The specific experimental manipulations at each stage included: introducing distraction during encoding, targeted memory reactivation during sleep-related consolidation, and selective spatial retrocues during retrieval. The theoretical motivations and scientific approach for the studies in each empirical chapter are outlined below.

At encoding: Distraction

Chapter II explores how manipulating attention during the encoding phase influences resulting memory-guided attention and memory quality. Attentional selection during the learning phase was influenced by including a highly salient distractor within the spatial learning task. The chapter explores the possibility of negative attentional guidance based on learnt information about distractors: Can we learn to ignore irrelevant, distracting information? Can we integrate knowledge about such distraction into our memory to form inhibitory tags on the relevant spatial location? Would this allow proactive suppression of the distractor location? And how would these putatively inhibitory processes affect the overall associated memory and its utility for subsequently guiding attention?

At consolidation: Targeted memory retrieval

Chapter III discusses a sleep study designed to evaluate the effects of sleep and sleep-based replay on memory and memory-guided attention. Using auditory stimuli

learnt in association with the visual search task within particular scene categories, I attempted targeted memory reactivation during a nap period, aiming to influence the consolidation process. The chapter encompasses a number of scientific enquiries: Are the memories employed in memory-guided attention subject to sleep-dependent consolidation? Can replay of the auditory context, in this case sounds that were not conceptually linked to the memory in question, be used to prime memories for selective consolidation? And can sleep-based consolidation affect the utility of memories for subsequently guiding attention?

At retrieval: Retrocues

Chapter IV provides, to my knowledge, the first evidence for effective retrocueing in long-term memory. The research is informed by the idea of selective retrieval, allowing internally guided attention to highlight memories and memory components for preferential recall. Over three studies, I observed that spatially informative retrocues were effective within long-term memory, allowing speeded access to mnemonic information, when contrasted with uninformative or misleading retrocues. Retrocueing benefits to memory precision were also observed in contrast to misleading retrocues. The chapter includes a discussion on the robustness, versatility, and utility of this novel experimental manipulation in assessing a variety of scientific pursuits.

In the light of the literature

Chapter V concludes the thesis by discussing the experimental findings in the light of the memory-attention literature. What evidence came to light for attentional influences at the encoding, consolidation, and retrieval of memories? Did lasting

impacts persist to affect the quality and utility of the memory? How do these observations fit within and complement the existing literature? Could the current set of results inform and inspire new research hypotheses? Some space is also given to contemplating problems encountered throughout the doctoral work and potential solutions are considered. I conclude with an outlook on how this doctoral work contributed to our understanding of the mutual influences between memory and attention.

Chapter II. Distraction during encoding

Introduction

Almost every moment, we are confronted with a vast breadth of perceptual information, too rich and manifold to be processed all at once. Instead, we have to prioritise information relevant for completing our goals and ignore the irrelevant aspects, regardless of how distracting they might be. At the same time, we need to be able to react to important external information, might it be an unexpected opportunity or a threat.

Traditionally, our visual attention has been seen to achieve this "balance" by a dichotomy of biases: Bottom-up and top-down attention. Bottom-up biases are driven by the physical salience of an event in the environment, which "captures" attention. This attention capture is often described as an automatic process (Neumann, 1984), or as a physiological reflex (exogenous attention, Posner, 1980). Top-down attention is the volitional act of directing attention according to our goals (Theeuwes, 2018). Our visual attention has hence been interpreted to be guided by a combination of these two attentional biases, one automatic and one voluntary (Jonides, 1981).

More recent studies on memory-guided attention have forced us to broaden this theoretical construct: Studies have shown that attention can also be biased toward learnt locations of items in memory (Summerfield et al, 2006; Summerfield et al 2011; Stokes et al 2012; Rosen et al, 2014; Rosen et al, 2018), even if these memories are not or have never been task-relevant and the attentional bias could be counterproductive (Soto et al, 2005; Olivers et al, 2006; Fan & Turk-Browne, 2015; van Moorselaar et al, 2019). These observations show that dividing attention into voluntary top-down and automatic bottom-up biases may be too simplistic, since internally generated biases may also act

automatically. Indeed, the hypothesis that there are sources of attentional bias unaccounted for in the traditional dichotomy has found many supporters in the literature (Awh et al, 2012; Theeuwes 2018, 2019; Anderson, 2016; Chelazzi et al, 2013; Jiang, 2018; Todd & Manaligod, 2018). Most notably, Awh and colleagues propose to extend the top-down/bottom-up attentional bias dichotomy to include current goals ("top-down"), physical salience ("bottom-up"), and selection history. Selection history is comprised of multiple components that bias attention, such as inter-trial priming (Maljkovic & Nakayama, 1994; Theeuwes & Van der Burg, 2011), reward history (Anderson, 2016, for review), and statistical learning (contextual cueing, Chun & Jiang, 1998). Nobre and Stokes (2019) argue that this definition should also encompass a variety of mnemonic information as potential guides to attention, such as evidenced by studies of memory-guided attention (e.g., Summerfield et al, 2006; Fan & Turk-Browne, 2015).

In most examples studied, the attention bias is positive, meaning that the association of reward or memory gives an attentional boost to the feature or location in question. However, it is interesting to consider the possibility of negative biases, whereby biases can deprioritise and inhibit processing of selective events. The consideration of negative biases based on physical salience is counterintuitive. There is disagreement in the literature over whether an external distractor cue could be used to suppress a salient stimulus (Arita et al, 2012; but Beck & Hollingworth, 2015; see Chelazzi et al, 2019). Some studies using external cues to orient goal-based attention suggest that inhibition of irrelevant items occurs as an indirect consequence of prioritising relevant items rather than by a direct cue-related inhibition of the irrelevant items (e.g., Noonan et al, 2016). However, long-term memory-based biases based on selection history can in principle lead to negative as well as positive biases.

The ability to use a distractor cue would necessitate a form of proactive suppression, in anticipation of the distractor. This can be seen as separately from retrospective, or reactive, suppression in response to the discovery of a distractor. There are differing opinions in the literature on the possibility of proactive suppression: Some theories and their proponents argue that any form of salient distractor necessitates the capture of attention (e.g., Theeuwes, 2004), even if we can subsequently rapidly disengage from the source of distraction (e.g., rapid disengagement hypothesis, Theeuwes et al, 1999; Theeuwes, 2010). Other theories suggest, that depending on task aims we can modulate our spatial priority map before a distractor appears to suppress the incoming signal prior to attentional capture (e.g., signal suppression hypothesis, Sawaki & Luck, 2010), allowing for the possibility of proactive suppression. Measuring proactive suppression often followed the logic of tracing the absence of measures indicating attentional capture: Behaviourally, this might be measured in reduced delay in reaction time or altered sensitivity to stimuli within salient singleton distractors in a suppression condition (Gaspelin et al, 2015). However, it could be argued that these results can also be accounted for by rapidly diverting attention, thus not showing a separate phenomenon to retrospective suppression. To differentiate these possibilities further and to dissociate attentional and suppression processes (e.g., see Chelazzi et al, 2019), researchers turned to the time-sensitive features of electrophysiology by analysing event-related potentials in response to the stimulus (Hickey et al, 2009; Sawaki & Luck, 2010) or oscillatory frequencies prior to stimulus onset, notably in the alpha band (e.g., Foxe & Snyder, 2011; Slagter et al, 2016). However, the ongoing variety of theoretical accounts speaks to the complexity in measuring and differentiating proactive and retroactive kinds of suppression, with studies supporting either account (for recent opinions, see Chelazzi et al, 2019; Liesefeld & Müller, 2019;

Theeuwes, 2019). Further, proactive and retroactive suppression mechanisms might not exist separately, but work in concert, with studies demonstrating evidence for both the signal suppression hypothesis and rapid disengagement (Wang et al, 2019). A distractor cue, or further the future-oriented use of the memory of a distractor, might help dissociating potential proactive suppression from distractor-bound retrospective suppression. Any modulation of behavioural and neural behaviour in response to a memory cue would be due to the anticipation of a distractor, especially if the cue is not actually succeeded by a distractor. By definition, there should be no retrospective, or reactive, suppression if no distractor appears that could be reacted to. Any differences in anticipation of or evoked by the mnemonic association with the distractor should reflect thus proactive processes designed to deal with the distractor.

It should be noted that the disagreement in the literature on whether it is feasible for a participant to utilise a distractor cue to ignore successfully a distractor may also be an artefact of task design. Noonan and colleagues (2016) observed that whilst participants appear to be able to learn to ignore a feature value or spatial location if the distractor value remains the same over trials or rather within a block of trials, the distractor cue becomes ineffective if it varies on a trial-by-trial basis. Indeed, Gaspelin and Luck (2019) have argued that the disagreement in the literature on the possibility of proactive suppression is due to the following differences in experimental design: Blocked design experiments present evidence for the existence of proactive distractor suppression (Leber & Egeth, 2006; Maljkovic & Nakayama, 1994; Gaspelin et al, 2015; Gaspelin et al, 2017; Sawaki & Luck, 2010; Gaspelin & Luck, 2018a), whilst experiments with trial-by-trial differences in the distractors find only retroactive suppression (Beck et al, 2018; Moher & Egeth, 2012). What implications does this observation have with regard to the potential of memory-based distractor suppression?

Within a task design, selection history appears to enable proactive distraction suppression. It has to be dissociated whether this suppression is caused by the immediate repetition of exposures, or whether the progressively established memory for the distraction information can lead to a negative bias.

Investigating whether memories for distracting information can lead to inhibitory biases is a new and important avenue for research: Can we learn to ignore an irrelevant, distracting item actively to optimise our attentional resources? This research question can be reframed to fit within the integrative framework proposed by Awh and colleagues as follows: If the selection history for an item is constantly to be selected against or ignored, can we learn to inhibit this item stably over time? It is important to note that suppression or inhibition of an item, feature, or location is not meant as simply not attending to it (Chelazzi et al, 2019). Suppression is meant to go beyond the mere removal of attention (Hickey et al, 2009; Sawaki & Luck, 2010; Donohue et al, 2018), suggesting that an object or feature can be in one of three states: attended, unattended, and suppressed.

There are studies showing that participants can learn over time to ignore and actively supress distractors under very specific circumstances (see Gaspelin & Luck, 2019, for a review on animal, and human behavioural and electrophysiological studies). Ipata and colleagues (2006) observed that, over repeated trials, macaque monkeys learnt to ignore salient distractors to the extent that their gaze was shifted away from the distractors and the neural representation of the distractor in the lateral intra-parietal (LIP) neurons was actively suppressed. Similarly, the expectation of a distractor appearing is enough to bias gaze away from the expected distractor location (Van der Stigchel & Theeuwes, 2006). Multiple studies corroborate the idea that distractors, such as colour singletons, can be avoided with increasing experience of the specific feature

value of the singleton (Gaspelin et al, 2017; Vatterott & Vecera, 2012; Cunningham & Egeth, 2016). Similar results were also reported for electrophysiological studies (Sawaki & Luck, 2010). There is thus evidence that inhibitory processes exist as the result of repeated exposure to the to-be-ignored value, regardless of whether they are the result of top-down processes (signal suppression hypothesis, Gaspelin & Luck, 2018b) or statistical learning (Theeuwes, 2018).

Most of the reported evidence for distractor suppression stems from visual search tasks, in which the response to distracting stimuli was attenuated or suppressed over time. Behavioural results in this context could result from some form of habituation that allows faster disengagement from stimuli retroactively or could point to mechanisms that enable proactive suppression of distraction if we know that it is about to happen. This makes it very difficult for pure visual search studies to dissociate between suppression theories and specifically provide evidence for proactive suppression, or at least engagement with distractor information.

Further, recent studies on distractor suppression have been using visual search tasks like the additional singleton paradigm (Theeuwes, 1990): The distractor, a salient singleton, appears as part of the search array, captures the attention of the observer and hence slows search for the actual target. Whilst these paradigms have been very influential in the field, they are limited in multiple ways. Firstly, the nature of the distractor itself has an essential influence on its effect (see Liesefeld & Müller, 2019), which makes it "unclear if all types of physically salient stimuli (especially sudden onsets and other dynamic stimuli) can be suppressed" (Gaspelin & Luck, 2019). Dynamic distractors with sudden onset however might be the most relevant to our cognition in daily life. In order to evaluate the effect of distractor suppression in a more ecologically valid setting, one would require varying kinds of distractors in more

natural task settings. Secondly, the effects of distractor suppression are influenced by the task setup: In the recent research on distractor studies, any inhibition or suppression was measured at the same time as the learning about the distractor occurred. This has the advantage that the development of suppression can be traced over time. However, it has the disadvantage that any suppression effect might be a mixture of progressively acquired suppression and active suppression in response to the appearing distractor. In other words, retrieving the information about the distractor suppression might be altered by the fact that a distractor is still present (cf. Wais et al, 2010). It would thus be interesting to tease these effects apart by measuring potential suppression due to distractor association when the distractor is merely remembered. A prerequisite for the idea of separating suppression learning and measurement is that inhibition cannot only be acquired but can also last over time (Turatto & Pascucci, 2016; Turatto et al, 2018).

Considering potential limitations of the current distractor paradigms introduced a new set of experimental questions: If we learn to suppress salient but irrelevant stimuli, do we learn to inhibit the associated features (such as spatial location) directly in a given context? Would such an "inhibitory tag" on the location exist only within the confine of a task, for example within a WM paradigm or in response to a distractor, or can this suppression remain stable over time and hence exist in LTM? Within the visual search-based tasks, target selection has been found to be less efficient if the target appeared in a distractor-associated location (Ferrante et al, 2018; Sauter et al, 2018; Wang & Theeuwes, 2018a). Indeed, Ferrante and colleagues found that increasing frequency of the distractor in a location was associated with decreasing attentional capture of a distractor and less efficient target selection, if the target was presented at a high-frequency distractor location. This evidence is consistent with the idea of an inhibitory tag, leading to suppression at the learnt location rather than just enhanced

attention for target locations. Ferrante and colleagues argue that the change in attentional biases is due to the statistical learning of target and distractor probability distributions, which would affect a proposed spatial priority map. Gaspelin and Luck (2018b) on the other hand argue that distractor suppression might be guided by pre-attentive feature information, and thus be processed before such a putative spatial priority map. The hypothesis that a potential inhibitory tag might be enacted proactively before any shift of attention has implications for the underlying neural mechanisms and the measures that might trace such an inhibitory tag.

Should such suppression mechanisms exist and last over a prolonged time in form of an associated inhibitory tag, how would we be able to measure it? If distractor association is processed pre-stimulus, one might expect modulation at an early sensory level, such as the P1 and N1. The early visual-evoked components P1 and N1 are of interest since they are known to be modulated, or amplified, by attention (Luck & Kappenman, 2012, ERP Components and Selective attention; see also Luck, 2005). Both P1 and N1 amplitudes are proposed to be larger for attended stimuli (sensory gain hypothesis, Hillyard et al, 1998). In a contrasting model, the P1 component has been suggested to be related to top-down inhibition (inhibition timing hypothesis, Klimesch, 2011) and thus have a role in blocking irrelevant sensory processing, showing larger amplitudes for inhibited stimuli (Slagter et al, 2016). Depending on theory, either expected, attended locations should be associated with the largest P1 amplitude (sensory gain hypothesis), or alternatively, distractor locations, followed by other unexpected, and then expected locations (inhibition timing hypothesis). Attentional effects on the P1 and N1 are also modulated by mnemonic effects: Positive attentional biases for target detection by memory (Summerfield et al, 2011) or reward association (Doallo et al, 2013) have resulted in an enhanced P1 component relative to a neutral

condition. Further, the same studies found some modulation of N1 amplitude and latency depending on condition. If distractor suppression occurs proactively, inhibiting the neural response to certain spatial locations, one might expect modulations in P1 and N1 amplitudes in response to targets in a distractor location, differing depending on the discussed hypotheses.

However, there is some evidence that suppression might not be processed during early visual potentials though, but rather at later stages in perception (Daffner et al, 2012). Some suppression studies have studied the N2pc as an index for whether a distractor captured attention (Sawaki & Luck, 2010). The N2pc is a negative deflection contralateral to the attended object beginning around 150-200ms after stimulus onset and has been associated with the covert deployment of visual attention (Gaspelin & Luck, 2018b). The N2pc has also been found to be attenuated in response to targets appearing in the memory cued location (Doallo et al, 2013). Within a task-set up that includes memory association for target and distractor locations, any modulation of the N2pc could be the result of either or both of the modulations and is thus unsuitable as a measure. Further, it is worth noting that suppression and attention do not necessarily share the same mechanisms and neural dynamics (Chelazzi et al, 2019).

More recent studies using visual search suppression paradigms have since focused on analysing the distractor positivity (Pd) as a measure of (reactive) suppression of the cortical response to a distractor (Hickey et al, 2009; Sawaki & Luck, 2010). The Pd component has been found to appear around 200-300 ms after stimulus onset (Chelazzi et al, 2019) as a response to distractors that fail to capture attention (Gaspar & McDonald, 2014; McDonald et al, 2013; Kiss et al, 2012). Indeed, the Pd fails to be elicited if the distractor is not salient enough (Chelazzi et al, 2019), which suggests that it might not be a suitable measure of distractor suppression if the distractor

is not present. It remains to be seen whether distractor association might be processed during a similar time window. A relevant component for this time window might be the positive inflection at around 200 ms post-stimulus presentation: The P2 component has been associated with higher-order attentional processing (Luck & Hillyard, 1994) and comparing perceptual input to expectation depending on context (Federmeier et al, 2005), showing an enhanced profile for repeated or memorised in contrast to novel visual stimuli (Curran & Dien, 2003).

Other electrophysiological measures associated with inhibition in the literature are the ERP components N2 and P3 in go/no go tasks (Falkenstein et al, 1999; Harper et al, 2014), or alpha-band oscillations, which increase in power contralaterally to distractor locations (Foxe & Snyder, 2011; Jensen & Mazaheri, 2010; Kelly et al, 2006). The negative inflection 200 ms post stimulus presentation (N2 component) is thought to be evoked in response to a visual stimulus deviating from context with subcomponents specifically being elicited by a template mismatch (Patel & Azzam, 2005). The P3 component (positive inflection around 300 ms post stimulus) has been interpreted as reflecting updating processes of memory, for example in response to stimuli that did not match the maintained memory (Patel & Azzam, 2005). The multitude of measures for potential inhibition or suppression is an indicator in itself for the ubiquity of the phenomenon and the complexity of measuring it. Indeed, Chelazzi and colleagues (2019) argue that depending on the nature of the distractor suppression multiple mechanisms might be engaged, which in turn would influence how such effects could be measured. This suggests that the appropriate measures will be dictated by the aims and layout of the respective experimental task.

In order to investigate the possibility of proactive suppression, it may be useful to move beyond using just visual search tasks. The memory-guided attention task

provides a good experimental approach in which it is possible to separate memory-based anticipation and modulatory consequences from the ongoing learning, selection, and inhibition processes involved in identifying a target within a visual search display (see Nobre & Mesulam, 2014; Nobre, 2018). The present study therefore adapted the memory-guided attention task to test whether inhibiting a salient distractor in a context results in learning an enduring inhibitory tag that can be deployed proactively to suppress anticipated distraction. The memory-guided attention paradigm (Summerfield et al, 2006) includes a separate learning task, a memory-cueing task, and a memory recall task. Within the Learning task, participants searched naturalistic scenes for a target key, whilst being distracted by a salient distractor in half of the trials. The aim was that a stable memory for the target and potential suppression of the distractor would develop after multiple exposures. The memory-guided attention task used the learnt scenes as a memory cue to guide attention to learnt target locations. The targets would appear very briefly, either in the expected, an unexpected or a distractor-associated location. I analysed attention, memory, and perceptual processes with behavioural and electrophysiological measures using electroencephalography (EEG). Primary electrophysiological analyses were focused on the early visual processing (P1 and N1 components) and exploratory analyses on later processing stages (P2, N2 and P3 components), in time windows known to be relevant for distractor suppression.

Method

Participants

Thirty-seven healthy and right-handed volunteers participated in the study. All participants had normal or corrected vision and reported no difficulties with colour

vision. Participants reported no history of neurological or psychiatric disorders and were on no medication that might affect their cognition. In total, data of 12 participants had to be excluded. Most of the exclusions occurred because participants had low performance on the memory task. Nine participants had memory-recall failure in more than 50% of trials. Data from the remaining individuals were excluded because of technical difficulties (1), hairstyle incompatibility with the EEG recording (1), and one participant withdrawing from the experiment before completion.

Consequently, the data of 25 participants were analysed (age range: 20 to 32; mean age = 23.9; *SD* = 3.2; 13 female). The study was approved by the University of Oxford Central University Research Ethics Committee (CUREC) and conducted in accordance with the Declaration of Helsinki. All participants gave written informed consent before testing and were financially reimbursed for their time (£40).

Materials

Two hundred and twenty-eight naturalistic scenes were used in the experiment. Stimuli were similar to those used by Summerfield and colleagues (2006), consisting of photographs of different indoor or outdoor views. These scenes were prepared using Matlab (Mathworks 20XXa/b, Natick, MA), and subtended 20° × 15° of visual angle (1000 × 750 pixels) at a viewing distance of 80 cm. For the Learning task, a small gold-coloured key (15 × 29 pixels; 0.29º × 0.57º visual angles) was placed in one of the four quadrants of the scene. Keys were preferably placed in a hidden location to encourage search and memory formation. For each scene, five different possibilities were generated: In one version the key was absent (used in Memory-Guided Attention and Memory tasks, *Figure 1b.-c.*) and the key was placed in a different quadrant in each of the remaining versions (Learning task, *Figure 1a.*). The various target locations were

used for counterbalancing purposes, ensuring that none of the effects occurred because of specific salience of target locations in particular scenes.

In addition to the target key, half of the scenes presented during the Learning task also contained a distracting stimulus. The distractor scenes remained the same across participants. The distractor consisted of a cartoon bee (98 × 98 pixels; 1.93° × 1.93° visual angles), which flashed on and off at a frequency of 10-11 Hz. The distractor was always placed in a different quadrant to the key. The relative position of the distractor to the target was counterbalanced in scenes and for participants.

The key in the Memory-guided Attention task appeared for a short exposure duration (100ms): The size of the key was increased (size: 25 × 49 pixels; visual angle 0.49° × 0.97°) and it appeared on a grey screen (RGB 128 128 128), following the scene stimulus without key (100ms) and a grey ISI of 1200 ms. No distractors were present at any stage in this task.

Eye-gaze position was tracked during the Learning and Memory tasks using a remote, infrared eye-tracking system (Eyelink 1000). A chin rest was used to stabilise head position and facilitate eye tracking. The height was adjusted for each participant. This set-up was kept the same for all tasks. EEG was recorded during performance of the Memory-guided Attention task.

Presentation software (version 14.3; Neurobehavioral Systems) was used for all aspects of experimental control, including stimulus presentation, recording behavioural responses and synchronizing experimental timing with the EEG and eyetracking recordings. The tasks were presented on a 22-inch Samsung SyncMaster 2233 monitor (resolution 1680 x 1050 pixels; screen width: 47 cm), using a refresh rate of 60Hz.

All statistical analyses were run with SPSS (Version 23.0.0.0) or R. For each statistical test, the respective assumptions were checked and a correction for multiple

comparison (Bonferroni) was applied where appropriate. Non-parametric tests were applied when appropriate. The significance level was set to $p < 0.05$ for all tests.

Design

In the experiment, participants performed three tasks: Learning, Memory-guided Attention, and Memory. The tasks are illustrated in *Figure 1*a-c.

Learning task

The Learning task was designed to establish new long-term memories for the specific spatial locations of a hidden key within a large set of naturalistic scenes, *Figure 1a*. Participants viewed 228 scenes containing a small hidden key in a randomized order.

Participants were instructed to search through each scene to find the hidden key. Half of the scenes contained a salient flashing cartoon bee as a distractor. Participants were not informed about the presence or purpose of the bee. Their task was to search the scenes to locate the target key. Maximal search time was set at 120 s. They were free to move their eyes and gaze position was recorded. Once the key was located, participants should click once with the mouse to activate a cursor. After the cursor appeared participants should click on the location of the key. The maximal time to move the cursor to the key location was set at 10 s. Participants were given feedback on screen (1000 ms) as to whether they had correctly identified the location of the key. Positive feedback "Target found" was given whenever the response was within 50 pixels of the key location, as measured in Euclidean distance. The next scene was either presented after a response was given by the participant or after the available search time expired (120 s). The inter-trial interval (ITI) was randomly drawn between 1000 and 2000 ms.

Feedback was also provided at the end of each of the seven learning blocks notifying the participants how many keys they had found during that block. The key location within the scene and the distractor association were held constant across learning blocks within a participant, but key locations differed across participants. The task design allowed for a within-subject analysis of the effect of learning block (7 levels) and distractor association (2 levels). Dependent measures included identification accuracy and reaction time.

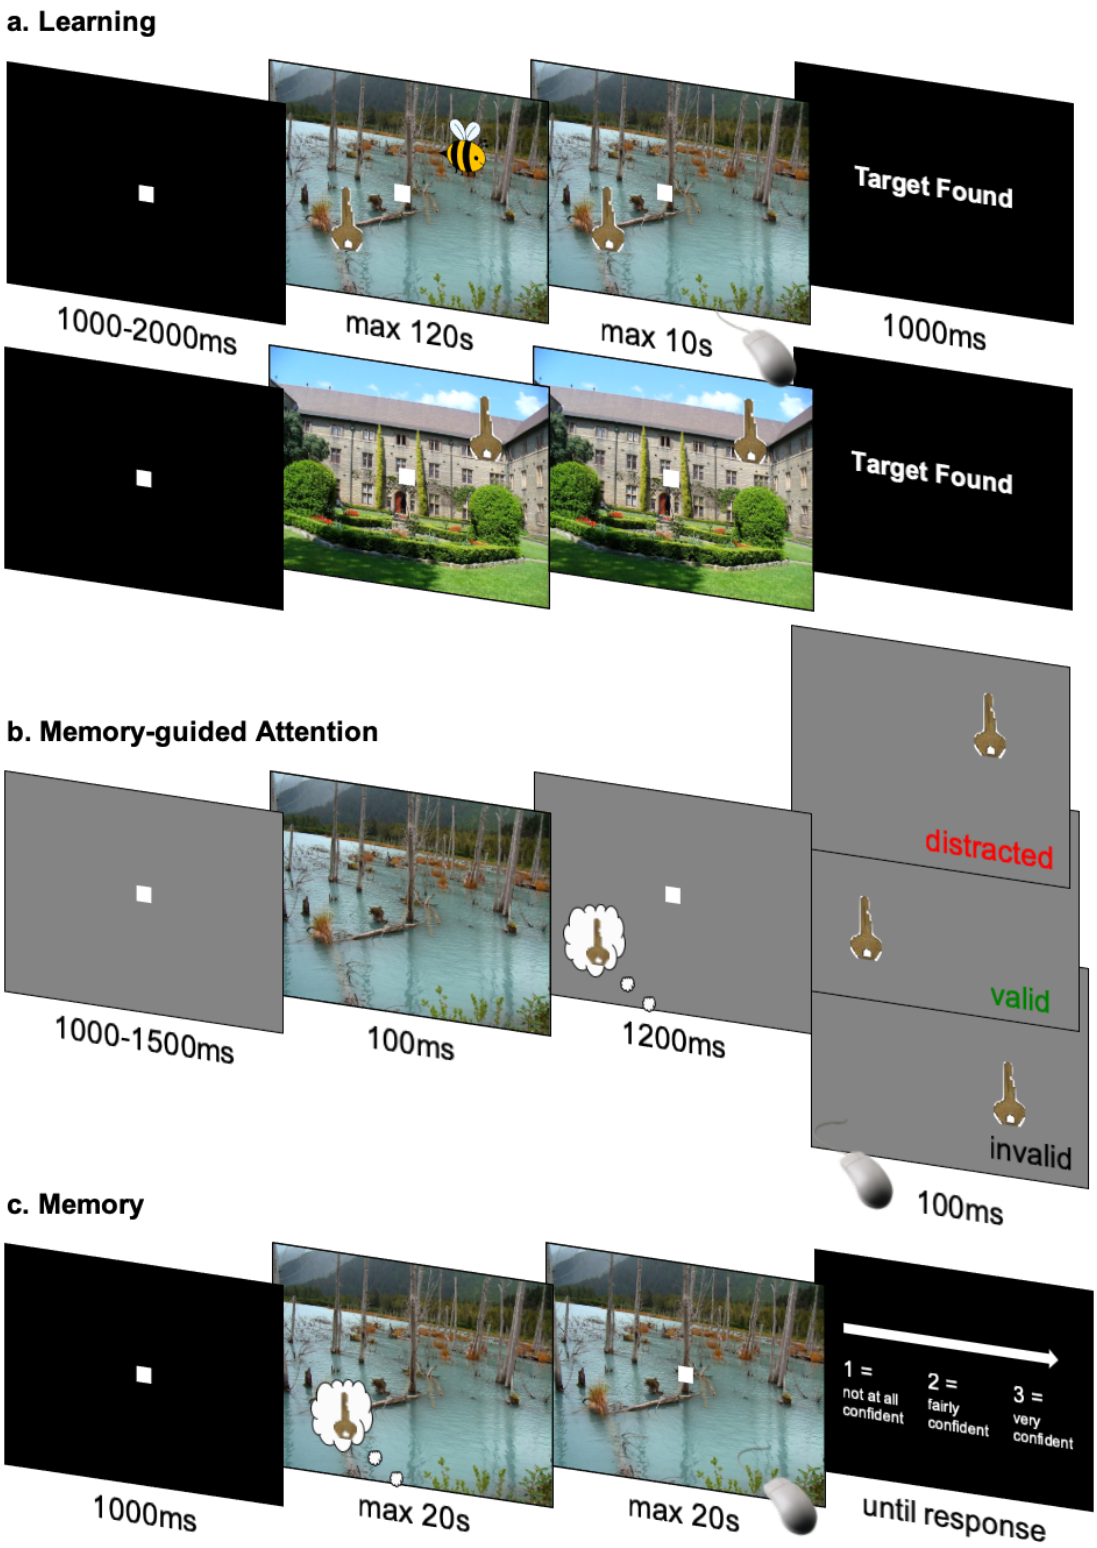


Figure 1. Schematic overview of the three tasks. Panel a. illustrates the Learning task: Participants were asked to search for the key in a complex scene. For half of the scenes a distracting bee was presented flashing at 10-11 Hz. Panel b. illustrates the Memory-guided Attention task: Each learnt scene was used as a memory cue to orient the attention towards the learnt key location. In 89 trials the key was presented in the learnt location, 91 in an invalid or 36 in the distracted location. The discrepancy is due to a coding error. Panel c. depicts the Memory task: Participants were asked to view the learnt scenes freely (without key or distractor present) and then indicate the spatial location of the key as remembered. Participants had to rate their confidence on a separate screen. Stimuli sizes are not to scale.

Memory-guided Attention task

The Memory-guided Attention task was a Posner-like cueing task in which memory for the spatial location of the target served as a spatial cue, *Figure 1b*.

In each trial, a previously learnt scene without embedded key appeared briefly on the screen (100 ms). After an interstimulus interval (ISI, 1200 ms) represented by a grey screen with fixation square, the target key flashed up quickly (100 ms). The key could appear in the learnt position (condition: valid), an unlearnt position (condition: invalid), or the location of the previous distractor (condition: distracted). Participants had to respond as accurately and as quickly as possible upon detecting the target by clicking on the mouse. Participants did not have to click on the target location and the target did not stay on screen. Instead, directly after the target was presented, the ITI began, again presenting a grey screen with a fixation square. The duration of the ITI was drawn from the range of 1000 – 1500 ms. Participants completed 228 trials, one containing each of the learnt scenes. Participants were required to maintain central fixation throughout but were given breaks at regular intervals (every twelve trials) to rest their eyes. Electroencephalography was recorded during the task. Triggers were sent for the appearance of a scene cue, the target, and any response by the participants.

Interspersed with target trials, participants were presented with 12 catch trials, in which no key appeared. These catch trials were drawn from the 228 learnt scenes and the 228 trials presented now included the 12 catch trials. Due to the fixed ISI, catch trials were useful to encourage participants to pay attention to the target key appearance, rather than to click automatically at the time targets were likely to appear. Catch trials were excluded from the main analysis but were evaluated separately to get a measure of response bias to non-existent targets.

The task was laid out according to the following logic. The scene served as a memory cue. The critical manipulation was the location where the target key occurred relative to its previous location in the Learning task and relative to the presence and location of the distracting stimulus during Learning. In the 'Valid' condition, the target appeared in its previously learnt location, in which case the memory validly cued its location. In the 'Invalid' condition, the target appeared at a different, unlearnt location. In this case the memory cue was invalid. Targets appearing at invalidly cued location could appear in any of the three quadrants not previously occupied by the target. In the 'Distracted' condition, the target appeared at the location previously occupied by the distracting bee. The distracted condition is a special case of the invalid-cue condition. For Valid and Invalid conditions, the learnt scene could be associated with the presence or absence of a distractor during learning.

The duration of 1200 ms for the ISI between memory cue and target presentation was chosen to investigate neural signatures as measured with EEG. The time following the onset of the memory cue will be denoted as cue period and the time following the onset of the target will be denoted as target period.

Regarding the cue period, the EEG data can be divided into two conditions: distractor-present or distractor-absent during the Learning task. For the target period, one can identify five conditions. The target key could appear in a previously learnt location with a distractor present (valid D+) or in a valid location but without distractor absent (valid D-). Alternatively, the target key could appear in an unlearnt, hence invalid, position with a distractor present (invalid D+) or with a distractor absent (invalid D-). These conditions allow a 2x2 comparison, contrasting the effects of memory-cue validity (valid/invalid) and distractor (presence/absence). Behavioural data were submitted to a 2 x 2 repeated-measures ANOVA accordingly.

In addition, the target could appear in the distractor position (distracted). This allows a 1x3 comparison for all distractor-associated scenes, namely comparing valid D+, invalid D+ and distracted, in a 1 x 3 repeated-measures ANOVA. All conditions were randomised in their temporal sequence across the Memory-guided Attention task.

Memory task

The Memory task was an explicit spatial recall task, *Figure 1c*. It was used to determine which scenes were associated with accurate memories for the target locations.

The design was almost identical to that of the Learning task. Participants were presented with the 228 learnt scenes in a randomized order. However, unlike the Learning task, there was no target stimulus or distractor present within the scene. The task of the participant was to indicate the location of the target associated with a specific scene in the Learning task.

Each scene appeared for a maximum duration of 20 s. Participants were instructed to click a mouse-button to activate the cursor once they remembered the target location. This activated the cursor, which participants then had to move to the remembered location and click again (max 20 s). If they were unsure, participants were asked to make their best guess for the target location, but to click in the centre of the screen if they had no recollection of the scene or target location at all. (All scenes with a centre click were excluded from the analysis of the Memory task later.) Further, participants had to rate the confidence of their responses using a three-point scale (1 = not at all confident, 2 = fairly confident, 3 = very confident). The confidence scale was presented on screen until a response was made. Trials were separated by an ITI of 1000 ms with a central fixation square. As in the other tasks, participants kept a stable position using

a chin rest, but they were allowed to move their eyes freely to explore the scene. Eye gaze was tracked throughout the first scene presentation. The main recorded dependent measures were RT and spatial error in Euclidean distance.

Procedure

Participants performed three experimental tasks over two consecutive days. On the first day, participants were informed that the study would involve three different computerised tasks – two using eye tracking and one using EEG. After reading the information sheet, participants provided informed consent. They were given time to read a task description including schematics of the Learning task. The experimenter explained the task again, allowing for questions. Participants were instructed to find as many keys as possible and to memorize the exact location within the scene. Participants then sat in the experimental booth, where the chin rest was adjusted to a comfortable height. Participants performed six blocks of the Learning task. They were able to take breaks after each block of Learning. The Learning task on Day 1 lasted around 165 minutes including breaks.

On the second day, participants performed the seventh block of the Learning task under equivalent conditions. The total duration for the learning block on Day 2 was around 25 min. Participants were given about an hour break between the Learning task and the other tasks to ensure that the learnt scenes were in some form of long-term memory. During this break, the EEG system was set up. Particulars about the EEG set up can be found in section *EEG data acquisition*. Participants' skin was prepared with an alcohol patch and an abrasive gel. Mastoid references and EOGs were placed first, before placing the electrode cap on the participants. During the EEG set up, participants read the task description including schematics for the Memory-guided Attention task

and were given opportunity to ask questions about it. Participants were also shown the effects of certain physiological behaviours on the EEG traces, such as clenching of the jaws, grinding teeth, blinking and moving eyes. Participants were encouraged not to do those motions during the EEG recording. When all set up was completed, participants moved back into the testing booth.

Prior to the Memory-guided Attention task, participants performed a short practice session with 12 trials to familiarise them with the task and performance requirements. During the practice trials, novel and not previously learnt scenes were used (including three catch trials). Practice trials were not part of the analysis. The task requirements particular to EEG recordings were practised in the practice task: Participants were encouraged to fixate their eyes calmly on the fixation square and not to move them around the screen. They were also encouraged to blink after giving a response in order not to interfere with the EEG recording for each trial. To avoid visual stress, the task included ample breaks for the participants to relax their eyes. The Memory-guided Attention Task was completed, taking around 20 minutes including breaks.

Once the Memory-guided Attention task was completed, the EEG set was removed. Participants completed the final task: The Memory task took around 20 minutes to complete. Subsequently, the participants were given the facilities to wash, condition, and dry their hair. Participants had the opportunity to be debriefed about the study aims. Lastly, they completed a payment form. The second day lasted around 135 minutes. In total the study a duration of five hours was allocated to the completion of this experiment over two days.

Behavioural Analysis

Learning task

The encoding of key location within each scene, from here on called target-context associations, was measured by the accuracy and reaction time (RT) to identify the target location over the course of the learning blocks. For statistical analyses, accuracy data were transformed using arcsine, and RT data using log. As accuracy data did not meet the assumptions of normality, I used equivalent Friedman tests over blocks for each of the distraction factors. For the analysis of the RT data, I conducted a repeated-measures ANOVA to test for effects of block (7 levels) and distraction association (2 levels: present, absent). To capture progressive learning, a linear contrast of the block factor was used. Trials were excluded based on whether no response was made and on reaction times (RTs) faster or slower than 3 standard deviations of the condition mean.

Memory-guided Attention task

The primary behavioural measure for the Memory-guided Attention task was RT, even though I recorded accuracy as well. For statistical analyses, RT data were log transformed. Accuracy data were merely reported as percentage of hits in target trials and false alarms in catch trials. RT data were analysed in two ways. The effects of validity and distractor association were measured in a 2 (valid, invalid) x 2 (distractor association: present, absent) repeated measures ANOVA. Further, distractor-associated trials only were contrasted in a 3-level repeated measures ANOVA (valid D+, invalid D+, distracted).

Trials in the Memory-guided Attention task were excluded 1) if they were catch trials in which no target appeared, 2) on the basis of the memory for the scene, and 3) on the basis of RTs faster or slower than 3 standard deviations of the condition mean. The second exclusion criterion is required as the scene is supposed to act as a memory cue towards certain spatial locations. If the participants retained no memory for the target-context association, the scene would not act as a memory cue. The exclusion was defined according to performance in the Memory task: Trials were excluded if participants had failed to make a response, or if the spatial error had been beyond 300 pixels away from the actual target location (in Euclidean distance). These exclusion criteria led to the exclusion of 25.70% of the total trials (*SE* = 1.87%) and the equivalent trials were excluded from the EEG analysis.

Memory task

The outcome measures of interest within the Memory task were RT and spatial error in Euclidean distance, as well as confidence ratings. For statistical analyses, RT and spatial error data were log transformed. Spatial errors from the participants' responses were compared with a chance level. Chance performance was estimated for each participant by using a permutation: Chance was defined as the average distance of hypothetical object response pairs. These were obtained by shuffling the x and y coordinates of actual object locations and responses over 1000 iterations, and calculating the Euclidean distance. The resulting measure mimicked the spatial error if responses were just given randomly.

The main conditions, whether the scene had been associated with a distractor during the Learning task or not, were contrasted using paired samples t-tests for both RT and spatial error data. This comparison for spatial errors was done over all

responses, and over accurate responses. Each scene was considered to be associated with an accurate memory if the participants located the target within a 300-pixel radius of its original location. Further, RT and spatial error data were analysed in accordance with associated confidence ratings: This was achieved by dividing the data along the three-point confidence scale, resulting in a 2 (distractor association: present, absent) x 3 (confidence rating: not at all, fairly, very) repeated measures ANOVA.

A supplemental analysis explored whether the distractor location distorted the spatial error for placing the target. A distortion measure was derived by calculating the difference in distance between the distractor location from the original target location and between the distractor location from the reported target location. No distortion relative to the distractor would result in a null distribution. A distortion towards the distractor would result in a positive distribution, and a distortion away from the distractor would result in a negative distribution. Trials were excluded from the Memory task analyses if 1) no response was made, or 2) the participant clicked in the centre of the screen to indicate they had no memory of having encountered this scene before.

EEG Data Acquisition

EEG data were collected with NuAmp amplifiers and Neuroscan data acquisition software (Compumedics Neuroscan, El Paso, TX) from 35 scalp sites, using Ag/AgCl electrodes mounted on an elastic cap, positioned according to the 10-20 international system (AEEGS, 1991). The reference electrodes were placed on the left and right mastoid bone. The left mastoid was used as an active reference during the experiment. Data were later re-referenced off-line to the algebraic average of the two reference electrodes. The common ground electrode was placed in the cap in the position of FpZ.

The montage included 7 midline electrode sites (FZ, FCZ, CZ, CPZ, PZ, POZ, and OZ) and 14 sites over each hemisphere (FP1/FP2, F7/F8, F3/F4, FT7/FT8, FC3/FC4, T7/T8, C3/C4, TP7/TP8, CP3/CP4, P7/P8, P3/P4, PO7/PO8, PO3/PO4 and O1/O2). In addition, vertical and horizontal Electrooculograms (EOG) were recorded. The vertical EOG was used for detecting blinks and was derived by subtracting signals between FP2 and an electrode placed under the right eye. The horizontal EOG was used to detect saccades and was derived by subtracting signals from electrodes on the outer side of right and left eyes. All impedances were kept below 10 kΩ. Signals were recorded with a low-pass filter of 200 Hz and a high-pass filter of 0.1 and were digitized at 1000 Hz.

EEG Analysis

Data pre-processing was performed offline using EEGLAB (version 14_1_2b, Delorme and Makeig, 2004) and customized Matlab scripts. EEG raw data were filtered offline between 1 and 40 Hz (finite impulse filter; filter order high-pass filter: 500; filter order low-pass filter: 100; for parameters chosen and ICA procedure see also Stropahl et al., 2018). Data were down-sampled to 500 Hz and segmented into consecutive 1-s time arbitrary/non-task related intervals. For two participants, signals from excessively noisy EEG electrodes were interpolated using the standard spherical method (Perrin et al, 1989; Ferree, 2000). Segments containing improbable data were rejected (*pop_jointprop*; *SD* = 3; M_{trials} = 120, SD = 32 of a total of M = 1388, SD = 107). The remaining data were submitted to extended independent-component analysis (ICA; 39 components) and the resulting unmixing weights were applied to the continuous raw data. Using the 1-s intervals, ICA components reflecting eye blinks, heartbeat, and other sources of non-cerebral activity were identified by visual inspection and removed from the actual data. This procedure resulted in the removal of 1–5 components per data set

(M = 3.28, SD = 1.02). Components identified as lateral eye movements were not excluded at this stage, but trials in which saccades occurred were removed after epoching in order to avoid conflating effects related to eye movements and to orienting attention. The ICA-corrected data were then re-referenced to the algebraic average of the left and right mastoid and high-pass filtered at 0.5 Hz (finite impulse filter, filter order: 1000). For one participant, an amplifier artefact led to the removal of two trials before the pruning stage that were not picked up by the joint probability rejection.

To derive event-related potentials (ERPs), data were epoched according to the memory cue (scene) and the target (key): Cue epochs were defined as -1.2 s to 1.2 s relative to the cue onset. Target epochs were equally defined as -1.2 s to 1.2 s relative to target onset. The epoch sizes were defined for purposes of pre-processing and data cleaning, and not the entire window was analysed statistically. Epochs were baselined using a window of -100 ms to 0 ms relative to cue or target onset. For time-frequency analyses, the epochs were defined between -1.2 s to 2.2 s relative to cue onset.

Any epochs containing residual artefacts were removed. Epochs containing non-stereotypical artefacts were pruned from the data based on kurtosis and probability of the activity based on local (single channel) and global (across channels) activity distributions (SD = 4; pop_jointprob, pep_rejkurt). This led to the exclusion of around 5 cue trials and 0.4 target trials on average per condition. Trials containing saccades at relevant periods of the epochs were also removed: All epochs were inspected manually and removed if saccades were detected within the time window of interest, either during cue presentation or just before or during target presentation. Additionally, the numbers of TF epochs were matched for each condition to ensure equal trial numbers within participants. This was achieved by selecting a random sample of the required number

of trials for each condition. One further trial was removed from EEG analyses for one participant due to a technical error.

For subsequent statistical analysis, data were transformed to Fieldtrip format (http://www.ru.nl/fcdonders/fieldtrip; Donders Institute for Brain, Cognition and Behaviour). In order to analyse lateralised effects of attention independent of stimulus side, electrode locations were converted into a relative space. For cue epochs, locations were contralateral or ipsilateral to the learnt target location.

Time-domain analysis

I wanted to investigate how memory-guided attention modulated the target anticipation and processing. In addition, I was interested in whether distractor association influences this modulation.

During the cue period, distractor and no-distractor associated EEG traces were compared. Any differences within EEG signature evoked by the cues would be due to the presence or absence of a distractor association for the scene in question. The effects of these conditions were assessed using a non-parametric cluster-based permutation test (Maris & Oostenveld, 2007). The test is equivalent to a non-parametric t test, thus comparing two experimental conditions at a time using a Monte Carlo significance probability. This is calculated by splitting the data in arbitrary conditions (equivalent to the size of the two experimental conditions), and calculate test statistics on these arbitrary partitions, specifically the maximum of the summed t values at the cluster level[1]. This test statistic was calculated for an alpha threshold of .05 over 1000

[1] The cluster-based test statistics are calculated by comparing all relevant data samples, selecting all samples over a specified threshold (p set to .05) and clustering resulting samples according to temporal, spatial and spectral adjacency. The cluster-level statistics are taken from these clusters (sum

iterations. The significance probability was calculated by comparing the proportion in which the randomised test statistics were larger than actually observed ones. For the statistical comparisons, alpha was pre-defined as .025, the appropriate value for a two-tailed test. Within the cue period, the cluster-based statistics were run over the entire time window of cue onset to target onset (0-1200ms), and over all electrode locations.

Analysis of ERPs elicited by targets followed the logic of the behavioural analyses. As noted above, the non-parametric permutation tests only reliably allow comparisons between two conditions. Thus, the predictions of the 2 x 2 ANOVA (memory cue: valid or invalid; distractor during learning: present or absent) and of the 1 x 3 ANOVA (memory cue: valid, invalid, or distracted) were tested in a pairwise fashion. Cluster-based analyses were conducted equivalent to the tests described above. However, data samples for target windows were analysed over predefined ERP windows of interest only, so the data were averaged over these time windows to account for the shorter temporal duration of the samples. The selection of the time windows was based on identifying the time course of the potentials in the grand-averaged waveforms across all conditions and electrodes. The time windows were defined as follows: P1 [110-150 ms], N1 [150-190 ms], P2 [220-260 ms], N2 [260-290 ms] and P3 [310-390 ms]. These epochs include a 20ms shift in the EEG signature, as a photodiode test confirmed a 20 ms delay in screen response. The permutation test averaged the ERP data over time for each of these predefined component windows. The effects of experimental conditions were tested on these major potentials elicited in the task across all electrode locations. The target period analyses were thus conducted constrained by

---

of t values within each cluster), of which the maximum is the test statistic used to evaluate the experimental conditions.

time, but unconstrained over electrode positions. Again, the pre-defined alpha was set as .025.

A set of further exploratory analyses aimed to understand whether validity-related differences in components were due to averaging across participants, or whether these effects were driven by individual differences. For this analysis, the average ERP values for each of the later ERP components (P2, N2, P3) were calculated across the whole brain and the entire component time window. The difference between valid and invalid served as the validity measure, which was then correlated with the other difference measures and the respective behavioural, normalised validity effect. The analysis was separated in distractor-associated trials and no distractor-associated trials for both ERP and behavioural measures. For example, the normalised behavioural validity effect was calculated as $\frac{M_{valid}-M_{invalid}}{M_{valid}+M_{invalid}}$, separately for distractor and no distractor scenes. Correlations were assessed with the Spearman's rank correlation coefficient.

Frequency-domain analysis

The preparation of time-frequency data differed from the time-locked data in the following way. Once in Fieldtrip, the TF data were prepared using a multitaper approach with Hanning wavelets, 5 cycles per frequency (temporal smoothing), and 0.3 Hz frequency smoothing, for frequencies between 1 and 30 Hz. The TF data were not baselined. The same non-parametric cluster-based permutation test (Maris & Oostenveld, 2007) as described above was used for the time-frequency analyses. For this set of analyses, I did not average over time windows, yet instead I averaged over predefined frequency ranges (theta: 4-7 and alpha: 8-12).

Results

Learning target locations with and without distraction

During the Learning task, participants identified the hidden keys within 228 scenes over seven learning blocks. Over the course of the learning blocks, increasing numbers of targets were located with increasing speed (mean ± SEM over all conditions; Accuracy - Block 1: 99.61 ± .09%, Block 6: 99.88 ± .05%, second-day Block 7: 99.95 ± .03%; Reaction Time - Block 1: 3.1 ± 0.17 sec, Block 6: search times 1.2 ± 0.04 sec; second day: Block 7: search times 1.3 ± 0.04 sec). Friedman tests showed that the accuracy improved significantly for blocks containing distractors ($X^2 = 31.84$, $p < .001$), *Figure 2*. However due to ceiling effects in the no-distractor condition, there was no significant improvement ($X^2 = 3.46$, $p = .749$). The 2 x 7 ANOVA on RT data revealed a main effect (F (2.46, 58.95) = 197.46, $p < .001$, $\eta_p^2 = .89$) and linear contrast (F (1, 24) = 321.49, $p < .001$, $\eta_p^2 = .93$) of blocks, *Figure 2*. There was also a main effect of distraction (F (1, 24) = 19.01, $p < .001$, $\eta_p^2 = .44$). There was no interaction (F (2.98, 71.759) = .81, $p = .490$, $\eta_p^2 = .03$), suggesting an equivalent rate of improvement in both distractor and no-distractor conditions.

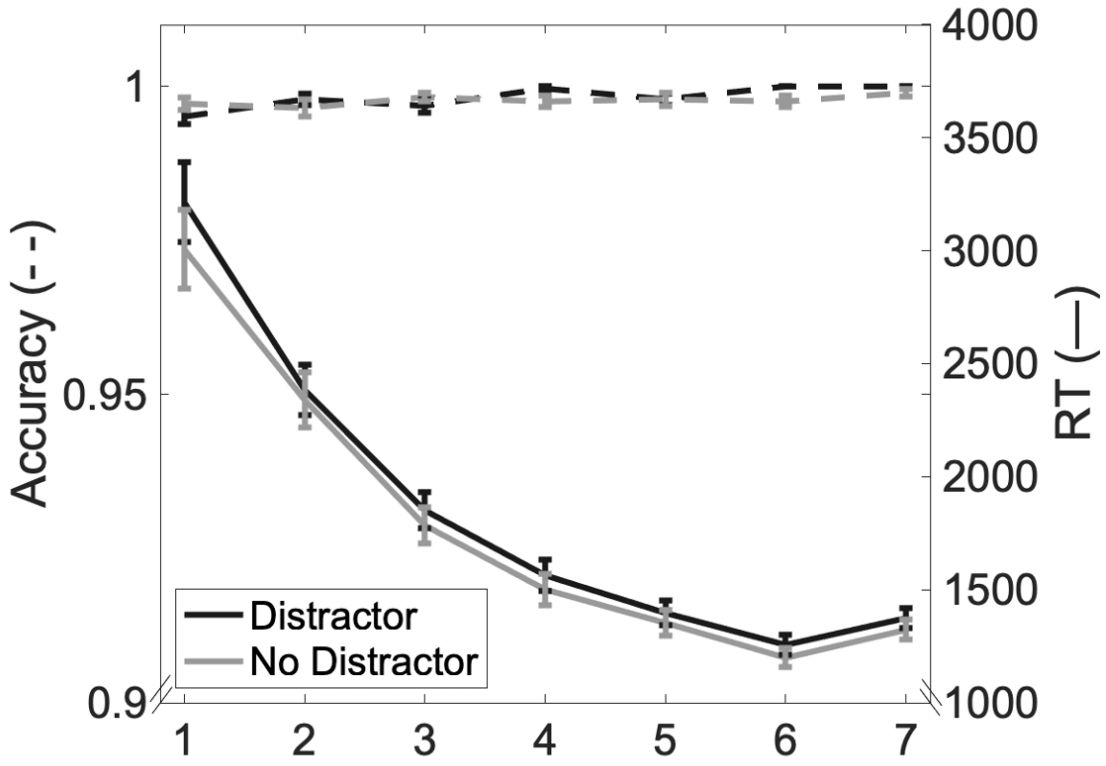


Figure 2 shows the increasing identification accuracy (dotted lines, in proportion of accurately chosen objects) and decreasing reaction times (continuous lines, in ms) over the three learning blocks. Dark grey depicts measures in response to scenes with distractor, light grey without distractor. Error bars are always standard error of the means.

Recalling target locations with and without prior distractor association

The explicit spatial Memory task confirmed that the participants retained accurate memories for target locations within at least some of the learnt scenes. Participants showed a mean spatial error of 170.39 pixels (*SE* = 13.36) for distractor scenes and 159.01 pixels (*SE* = 11.50) for no distractor scenes, equivalent to a mean hit rate of 80.39% (SE = 2.04%). Chance performance was estimated to be at 450.33 pixels (SE = 2.05). Responses to both distractor scenes and no distractor scenes were significantly above chance, t (24) = -11.30, $p < .001$, $d = 2.26$ and t (24) = -13.19, $p < .001$, $d = 2.64$ respectively. However, a paired-samples t-test showed no significant difference

between the spatial errors in response to distractor association, t (24) = 1.54, p = .135, d = .31, *Figure 3*. There was also no significant difference, if only responses regarded as hits were considered, t (24) = 1.16, p = .258, d = .23 (Distractor: M = 78.34, SE = 3.79; No distractor: M = 75.31, SE = 3.29). Reaction times to distractor (M = 1.69 s, SE = .15) and no distractor-associated scenes (M = 1.67 s, SE = .13) also did not differ significantly across all trials, t (24) = .13, p = .899, d = .03.

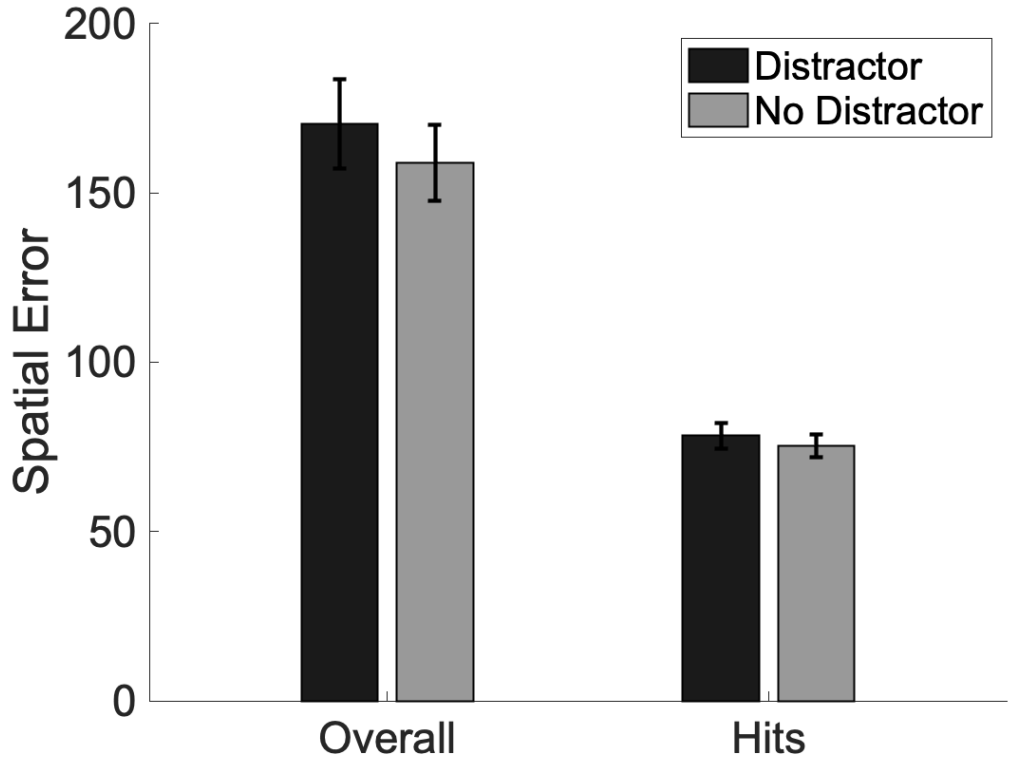


Figure 3. The spatial error in pixel distance is shown for the scenes associated with (dark grey) and without a distractor (light grey). Separate bar plots are shown for errors over all trials, and for trials considered as hits (within 300 pixels radius). Error bars are always standard error of the means.

Analysis of reaction times and localisation error as a function of associated distraction and confidence showed participants to have good awareness of the quality of their memory. Participants were responding significantly and progressively faster with increasing confidence (RT in s: *M*, *SE*; Confidence level 1: 1.92, .19; Confidence level 2: 1.88, .14; Confidence level 3: 1.46, .12), as evidenced by significant main effect

(F (1.36, 32.61) = 21.54, p < .001, η_p^2 = .47) and linear contrast (F (1, 24) = 21.12, p < .001, η_p^2 = .47) of confidence, *Figure 4a*. There was no main effect of distractor on RT (F (1, 24) = .70, p = .411, η_p^2 = .03), hence no significant difference between responses to scene with (M = 1.75 s, SE = .14) and without distractor association (M = 1.80 s, SE = .15). There was no significant interaction (F (1.53, 36.67) = 2.08, p = .149, η_p^2 = .08). The 2 x 3 repeated-measures ANOVA on spatial-error data revealed a similar pattern of results: Participants were increasingly more accurate with higher confidence (spatial error in pixels: M, SE; Confidence level 1: 329.19, 12.52; Confidence level 2: 180.13, 14.38; Confidence level 3: 59.39, 4.47), shown in a significant main effect (F (2, 48) = 350.28, p < .001, η_p^2 = .94) and linear contrast of confidence (F (1, 24) = 750.68, p < .001, η_p^2 = .97), *Figure 4b*. There was again no main effect of distractor (F (1, 24) = .04, p = .843, η_p^2 < .01), suggesting similar spatial errors in response to distractor (M = 188.74 pixels, SE = 10.84) and no-distractor associated scenes (M = 190.41 pixels, SE = 9.09. There was also no interaction (F (2, 48) = .61, p = .547, η_p^2 = .03).

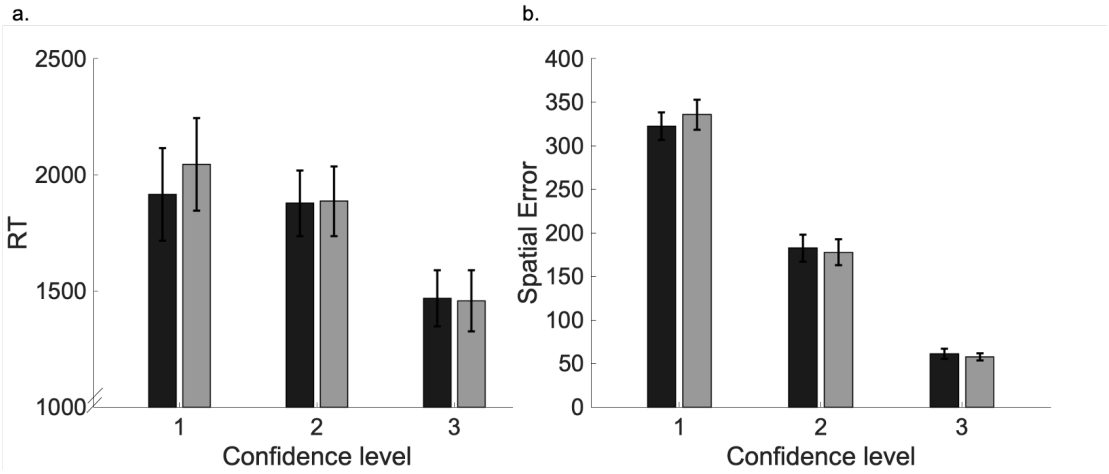


Figure 4. Measures of RT (panel a.) and spatial error (panel b.) are represented depending on associated levels of confidence (low, fairly, high). For each measure and level, the results are split into scenes with (dark grey) and without distractor association (light grey). Due to varying confidence ratings, the trial numbers differ for the different levels of confidence and distractor association (Not, Fairly, Very (M, SE); Distractor: 31.12, 2.71; 32.40, 2.09; 50.32, 3.29; No Distractor: 28.72, 2.58; 30.08, 1.78; 55.04, 2.99). Increasing confidence levels are associated with a decrease in RT and spatial error. There is no visible difference between distractor and no distractor scenes. Error bars represent the standard error of the mean.

Lingering distraction effects on established long-term memory

An analysis aimed at identifying systematic biases in recalling the previously learnt target location as a function of the associated distractor indicated that the distractor location exerted an attractor effect on the target recall reports. A one-sample t test showed that the localisation errors for reporting target locations were closer to the distractor location than would be expected by chance. Rather than being distributed around zero, the average difference in distance values from the previously learnt target vs. from the associated distracter were significantly larger than zero (M = 39.32, SE = 4.88; t (24) = 7.89, p < .001, d = 1.58) (*Figure 5*). This distortion remained reliable even when only considering highly accurate trials ('hits', errors within 300 pixels of the actual target location: M = 12.97, SE = 2.38; t (24) = 5.34, p < .001, d = 1.07).

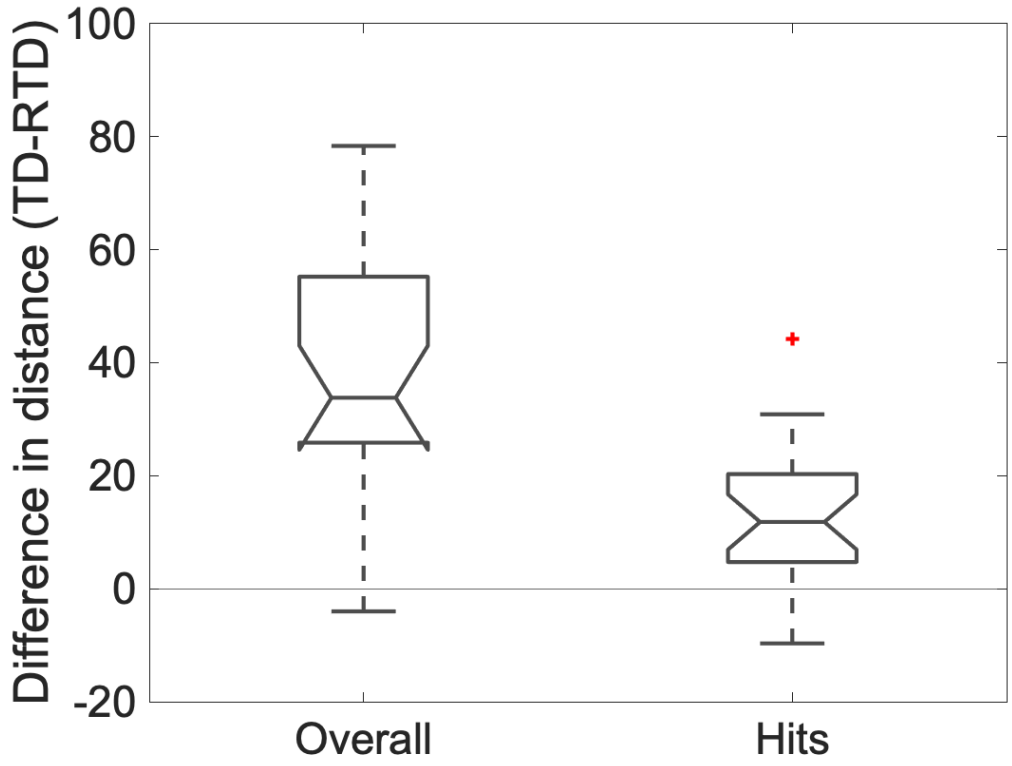


Figure 5. The bar charts show the difference in distance between target - distractor and reported target – distractor, both for all distractor scenes (overall) and just for distractor scenes in which target localisation was reported accurately (hits, within 300 pixels of the original location). The centre line of the bar plots

denotes the mean, and the edges of the boxes the respective quartiles. Values larger than zero reflect that the targets were reported closer to the distractor than their original location.

Memory-based attention with and without prior distractor association

During the Memory-guided Attention task, detection speed for targets flashing up after a memory scene cue was significantly influenced by the content of LTM. Overall, identification accuracy was at ceiling: Participants responded to 98.83% (*SE* = 0.24) of targets. On average, participants misidentified a target in 7.33% (*SE* = 1.72) of catch trials. Only accurate trials were used for the following analysis (*M* = 169.4, *SE* = 4.26).

A 2 x 2 ANOVA with the factor validity (valid/invalid) and distraction (distracted/non-distracted) was calculated for the reaction times. The main effect for validity was significant F (1, 24) = 60.65, $p < .001$, $\eta_p^2 = .72$, *Figure 6.* Responses after valid memory cues were faster (*M* = 348.39, *SE* = 9.44) than in invalid memory-cues condition (*M* = 373.62, *SE* = 11.02). No significant main effect for distraction was observed, F (1, 24) = 1.38, $p = .252$, $\eta_p^2 = .05$ (distractor: *M* = 362.34, *SE* = 10.09, no distractor: *M* = 359.67, *SE* = 10.28). There was also no differential effect of validity on distraction, as shown by a lack of interaction, F (1, 24) < .01, $p = .985$, $\eta_p^2 < .01$.

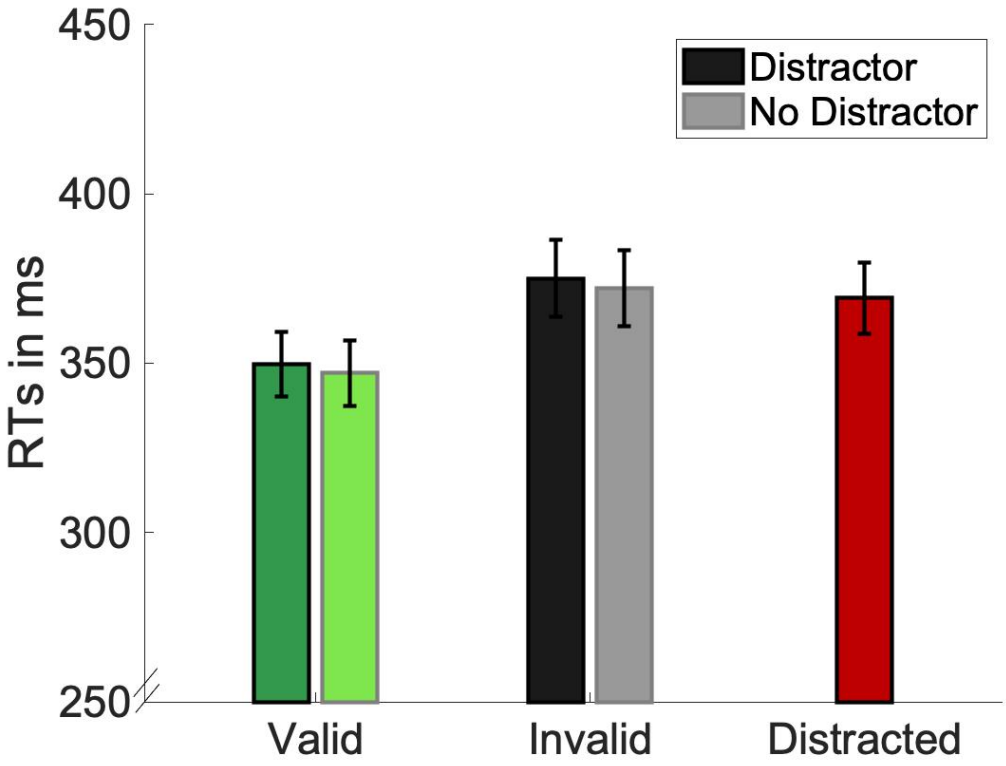


Figure 6. The graph displays reaction times in response to target appearing in the Memory-guided Attention task. Error bars represent standard errors of the mean. A dark grey outline around the bars reflects a distractor association of the scenes, whilst trials involving scenes without a distractor association are shown with a light grey outline. RTs in valid trials (green) appear to be shorter than in invalid trials (grey), in response to both distractor and no distractor associated scenes. Distractor association does not seem to affect RTs noticeably. Distracted trials (red) appear to be slower than valid trials, yet not different from invalid trials.

Comparison of the three types of trials associated with previous distraction, showed that response times for targets appearing at the previous distractor location was intermediate to that for targets appearing at validly and invalidly cued locations. A 1 x 3 ANOVA with the factor validity (valid/invalid/distracted location, log-transformed) yielded a main effect of location, F (2, 48) = 17.40, $p < .001$, $\eta_p^2 = .42$ (*Figure 6*).

Post-hoc t-tests showed that reaction times to valid locations (M = 349.63, SE = 9.52) are different from those to invalid locations (M = 375.05, SE = 11.29), t (24) = -5.13, $p < .001$, d = 1.03. The RTs to distracted locations (M = 369.22, SE = 10.62) were different from valid locations, t (24) = -4.95, $p < .001$, d = .99, but not from invalid, t

(24) = -1.19, p = .247, d = .24. However, there was also a significant linear contrast of the conditions, F (1, 24) = 26.28, p < .001, η_p^2 = .52.

Distraction association modulates the neural response to memory cues

Analysis of the cue-related neural responses focused on the question "*Was there an effect of previous distractor associations?"*, contrasting memory cues that had previously been associated with a distractor with ones that had not. A cluster-based comparison of ERPs over the entire cue-related epoch (0 to 1200 ms relative to cue onset) revealed a significant difference between cues previously associated with a distractor or with no distractor (p = .034, *Figure 7*). The distractor-associated memory cues showed a greater negative potential between 538 - 616ms post-cue. The effect was broadly distributed over frontal and central midline and right electrodes. No effects were observed in the cluster-based analyses of the induced time-frequency activity.

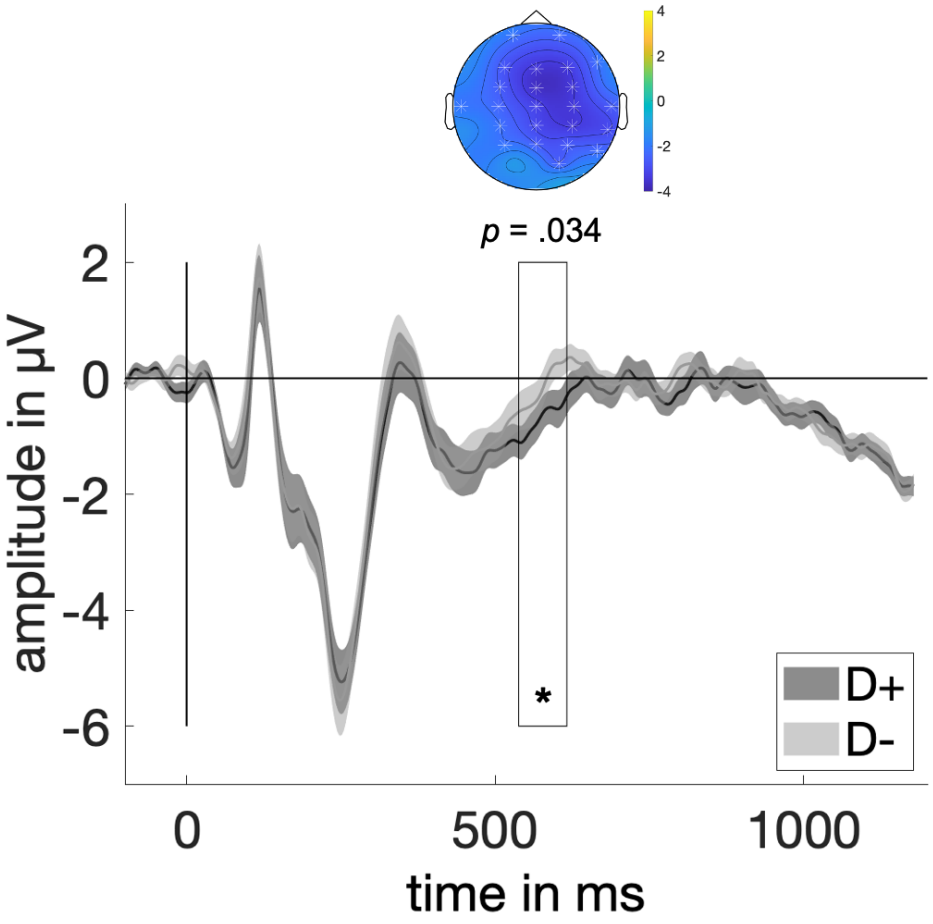


Figure 7. The ERP signal over all channels in response to the scene cue showed a late emerging cluster between distractor and no distractor scenes. The light grey line shows the ERP wave in response to no distractor scenes, whilst the dark grey line shows the ERP wave to distractor scenes. The shaded areas represent the standard error of the mean. The box denotes the time window during which the conditions show a significant difference, over sensors depicted in the topography.

Distraction association modulates effect of memory on target processing

The target-related ERPs were analysed in parallel to the behavioural analysis, namely attempting to answer the questions "*Was there an effect of previous distractor association?"* (Analysis 1) and "*Was there a differential treatment of the specific distractor location?"* (Analysis 2).

Effects of distractor association on target processing

Figure *8* shows the ERPs elicited by targets following valid and invalid cues separated according to whether a distractor had been previously associated with the scene acting as the memory cue. Differences were tested over all sensors within restricted latency windows in which individual potentials were most pronounced. No significant differences were observed in the amplitudes of the P1 or N1 potentials. However, the cluster-based permutation test revealed a significant difference in the P2 between the valid and invalid cueing condition (p = .013) for distractor trials (D+) only. In this latency range, the difference was most pronounced over fronto-central electrodes. The significant difference between the valid and invalid conditions in distractor trials also extended to the N2 latency (*Figure 8*). The two most pronounced differences appeared over ipsi- and contralateral posterior sensors, p = .033, and p = .035 respectively. No such difference was observed for no distractor trials (D-). For both the P2 and the N2 time windows the effects suggest a differential impact of prior distraction on the memory-based orienting effects. The cluster-based permutation tests also revealed significant differences between the valid and invalid conditions in the P3 for both distractor (p = .039) and no distractor trials (p = .039). In this latency range, there was an attenuation of the P3 potential elicited by validly cued targets over central-posterior sensors.

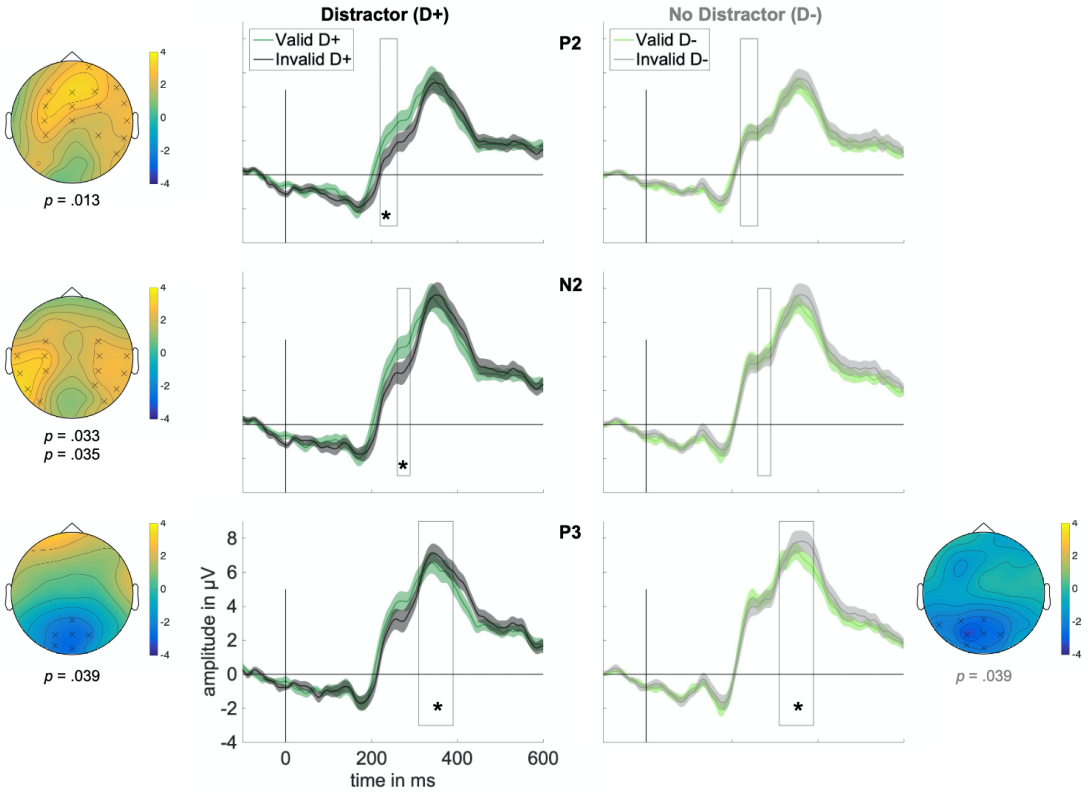


Figure 8. The green lines show the ERP waveforms in response to a target appearing in the valid location, whilst grey/black ones show the respective ERP signals in response to invalid target locations. The darker colours in the first column show trials associated with distractors, whilst the lighter colours in the second column show trials without distractors. The shaded areas denote the standard error of the mean, and the highlighted rectangles the latency of the tested differences. The topographies show the differences emerging in the cluster-based analyses, and these sensor clusters are used as a basis for the respective waveforms. If no difference emerged, the equivalent sensor cluster for the distractor trials was used. The first two rows of data show validity differences for distractor trials only in the P2 and N2 latencies, whilst the third shows significant validity effects for both distractor and no distractor trials in the P3 window.

Further exploratory analyses attempted to illuminate whether the emerging validity pattern for distractor scenes was an artefact of the ERP traces being averaged across participants, or whether indeed, for example, the participants showing the largest P2 components were also the ones with the strongest N2 component. Whilst keeping in mind the limitations of this analysis given the sample size consideration, it might help to understand whether the distractor-associated validity effects in the P2 and N2 windows are driven or rather masked by individual differences. The magnitude of the validity differences in distractor scenes was highly correlated across components (P2 vs N2: r (24) = .75, p < .001, P2 vs P3: r (24) = .67, p < .001, N2 vs P3: r (24) = .73, p

< .001). For no distractor scenes, this relationship was not as pronounced, showing only a significant correlation of the adjacent components P2 and N2 (P2 vs N2: r (24) = .71, p < .001, P2 vs P3: r (24) = .33, p = .107, N2 vs P3: r (24) = .73, p = .095). Interestingly, the validity-based differences in the component windows were not correlated with behavioural validity effect for distractor-associated (P2 vs behavioural: r (24) = -.27, p = .190, N2 vs behavioural: r (24) = -.27, p = .186, P3 vs behavioural: r (24) = -.37, p = .069) or no distractor associated scenes (P2 vs behavioural: r (24) = -.26, p = .207, N2 vs behavioural: r (24) = .12, p = .561, P3 vs behavioural: r (24) = .33, p = .112).

Was there a differential treatment of the specific distractor location?

Figure 9 also shows the target-related ERPs, expanding on the first (distractor-associated) column in *Figure 8*. The waveforms elicited by targets appearing at the location of the previous distractor (distracted) are plotted alongside those appearing at the previous target (valid) or at another (invalid) location. Cluster-based permutation tests showed no clusters in the P1 or N1 time windows (p not applicable).

Analysis in the P2 time window revealed a significant difference between the distracted and the invalid-cueing condition (p = .033) but not the valid-cueing condition (no cluster formed). The difference was most pronounced in fronto-central electrodes, similar to the electrodes in the valid D+ versus invalid D+ difference noted above. The pattern for the N2 reverses; ERPs in the distracted condition differ from those in the valid-cueing condition (p = .032) but not the invalid-cueing condition (no cluster formed). The contralateral posterior distribution of the effect is again similar to that observed in the valid D+ versus invalid D+ comparison above. During the P3 time window, ERPs in the distracted condition continued to differ from the valid-cueing (p

= .040), but not the invalid-cueing (no cluster formed) condition. The P3 modulation occurred over central posterior sensors, similar to the previous P3 contrasts above.

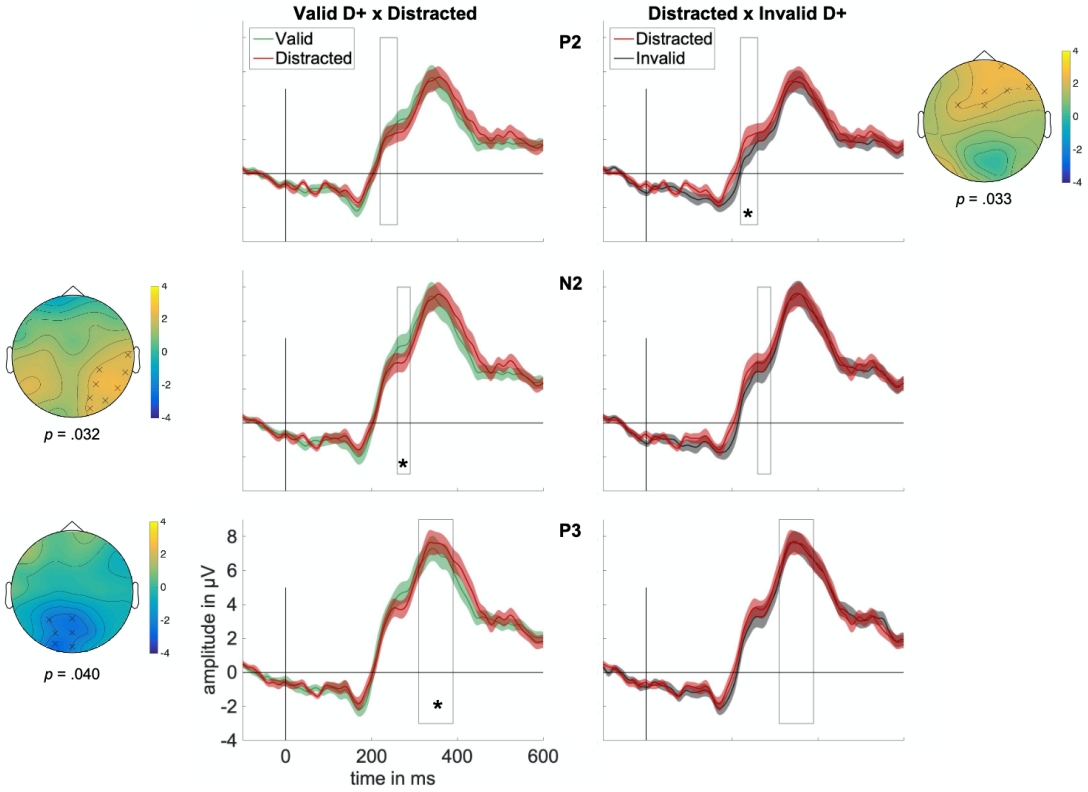


Figure 9. The red lines show the ERP waveforms in response to targets in the distracted condition. In the left column these are compared to validly cued targets in the distractor-associated condition (green lines) and in the right column these are compared to invalidly cued targets in the distractor-associated condition (grey lines). The shaded areas represent the standard error of the mean and the highlighted rectangles the windows of interest. The waveforms are averaged over the electrodes in showing significant effects in the time window. When no effect was present, the same electrodes are plotted as the effect in the other comparison for that time window.

Discussion

The task was successful in generating learning of relevant target associations in complex environments, as well as revealing that the irrelevant stimuli were distracting. However, the hypothesis that the learning of locations of distracting irrelevant stimuli might lead to an inhibitory tag was not unambiguously supported by the current findings.

In the Learning task, participants improved at locating the target over exposures and there is evidence that the distractor interfered with the visual search. The distraction effect did not prevent participants from locating the target or learning the target location progressively. If participants had been able to suppress the distractor fully with repeated exposures, one would have expected the size of the distraction effect to diminish. However, the distraction effect persisted over blocks of learning, suggesting that there were either insufficient exposures to generate an effective inhibitory tag, or that the features of the distractor (sudden onset, dynamic) were too salient to be inhibited. A potential follow-up study would focus on increased learning exposures to allow potential suppression to build up. This could be achieved via overtraining in a remote/online testing setup.

Typical measures of memory recall showed sufficient learning of target locations within the naturalistic scenes, and no effect of previous distraction on overall memory precision and quality. It is noteworthy though that the observed memory performance was worse than in previous, comparable studies. This might be an unintended consequence of distraction or due to the higher volume of scenes. Either way, the poorer memory performance might have contributed to the lack of a clear effect of distraction and could also benefit from more learning exposures. However, participants showed good meta-awareness of the quality of their memories.

Even though there was no effect of distraction on overall memory precision, more subtle analyses revealed remaining traces of the distractor association. Targets were reported systematically closer to the distractor, as if pulled towards it. This might suggest that the distractor continued to act, at least in some cases as a second attractor. Unfortunately, there was no measure for the memory of the distractor location, to ensure that the distractor continued to be seen as irrelevant for the task. This created

difficulties for assessing the full extent of participants distractor memories and their potential bias towards it.

Regardless of distractor association, the acquired memories acted as guides to orient attention towards the relevant target memories, in line with previous findings of memory-based attention. There was no behavioural evidence that the previous distractor location was suppressed. On the contrary, and in line with memory-recall biases, performance for detecting targets at previous distractor locations was intermediary. This might again suggest that the distractor location may have been more seen as a second target location by at least some participants and acted as an attractor to attentional biases. It is important to consider these results in the light that participants were unable to suppress the distractor during the Learning task, indicating that no inhibitory tag had been developed. The development of an attentional immunity to the distractor is obviously crucial to evaluate prolonged effects of such suppression. The results in the Memory-guided Attention task might have been more conclusive with a larger set of better trained associations. This could be achieved via remote testing.

Neural data recorded during the Memory-guided Attention task revealed only modest effects of memory-guided attention. I observed no early modulation of visual potentials, which would have been expected if attention had exerted a measurable, even proactive, influence on the initial perceptual processes. For example, if attention had emphasised the memorised, valid location, one should observe a higher amplitude of the P1 component for targets appearing the valid location and an inhibition of the P1 for targets in unlearnt locations (sensory gain hypothesis, Hillyard et al, 1998). Alternatively, should the P1 reflect top-down inhibition processes, one might expect the largest P1 amplitude for distracted stimulus locations, followed by invalid, and then valid (inhibition timing hypothesis, Klimesch, 2011). However, I did not observe any

of these hypothesised possibilities of inhibition of the P1 or N1 components. This could again be due to lack of the development of distractor suppression, and few trial numbers. Due to the nature of the task, only few trials per subject per condition remained to the detriment of the ERP analyses. Larger and over-trained data sets could be achieved via remote learning, with participants only being tested in person for the Memory-guided Attention task with electrophysiological recordings. Nevertheless, I observed later modulation of ERP components starting with the P2.

The exploratory analyses of the neural data revealed traces of distractor association and differential or more effortful target processing in the previous distractor location. There were cue-related effects visible broadly fronto-central, suggesting some processing of distractor association prior to target presentation. The target-related processing showed intermediate level processing of the distractor location. More specifically, I expected an enhanced P2 for targets in the learnt location, along-side the idea of the P2 component as representing a comparison between perceptual input and content-dependent expectation (Federmeier et al, 2005), enhanced for memorised targets (Curran & Dien, 2003). Interestingly, contrasting P2 amplitudes between learnt and unlearnt locations were only visible within distractor-associated contexts, as if this higher-order attentional processing (Luck & Hillyard, 1994) was sharpened by the distractor association. Further, the former distractor location itself also showed an enhanced P2 profile in comparison to "other unexpected" locations, supporting the idea of the distractor as a second relevant, or memorised item. With the subsequent N2 window, I expected stronger negative deflections for targets in unlearnt locations, demonstrating the response to a mismatch or deviation from context (Patel & Azzam, 2005). Again, the conditions only showed a significant difference in N2 deflections if a distractor had been associated with the scene, producing stronger N2 components for

unlearnt, unexpected locations. In relation to the former distractor location, it did not differ significantly from any other unlearnt location, yet it did differ significantly from the learnt, valid location. This pattern agrees with my hypotheses, as both targets in the distractor and any unlearnt location would have represented a mismatch to the learnt spatial template. The P3 component is thought to reflect updating processes (Patel & Azzam, 2005), incorporating novel, not memorised information. As expected, any target location that was not learnt during the learning phase resulted in a significantly elevated P3 amplitude, irrespective of distractor association or location. What can be inferred from this pattern of results in the light of the specific distractor location and the association of distractor context?

The P2 component, enhanced for target features or infrequent targets (Luck & Hillyard, 1994) was more pronounced for distracted locations in comparison to "neutral" unlearnt locations, which again suggests that it may have been seen as a second target to a certain extent. However, during later components N2 and P3, the response to distracted locations became similar to other invalid locations, and significantly different from the valid condition. The N2 component is described as denoting a deviation in form or context from a prevailing stimulus (in this case the target location), and the parietal P3 as a broad recognition or memory updating processes, maximal for mismatches (Patel & Azzam, 2005). Hence, whilst the distractor location might be seen as a relevant "attractor" location, it is still processed as a mismatch to the actual target location during later stages of processing. The distractor and its associated location might have perceived as a combination of distractor and second target. This could explain the absence of a behavioural effect differentiating target detection in distracted and invalid locations.

In line with these arguments, the mere presence of a distractor association appears to have amplified any effects of target related processing: Distractor presence during the memory acquisition resulted in significant P2 and N2 components. This may have been due to an amplified response to target matches for the valid condition (P2) and a stronger mismatch for invalid conditions (N2). The interpretation of these results as strengthened by the distractor context is supported by the findings of the exploratory analyses: Participants with strong validity effects in some components also showed stronger effects for other components, in particular for distractor-associated scenes. The distractor context appears to amplify these validity differences across conditions. Notably, this relationship appears to be irrespective of behavioural output, as the behavioural validity effect did not correlation with the neural measures of validity. The exaggeration of memory features due to distraction association is reminiscent of behavioural enhancement of memory encoding or retrieval in the presence of distractor suppression (Markant & Amso, 2013; Park et al, 2014; Markant et al 2015). The typical P3 for memory updating is seen for unexpected locations regardless of distractor association. The pattern of ERP results again stresses the importance of the P2 and N2 time windows for processing distracting information. This study predates the majority of the literature on the distractor positivity, but it is noteworthy that the relevant time window for processing distractor suppression matches the time windows in which ERP component effects were found. Despite differing task designs, the current findings are hence consistent with previous literature on the importance of the time window of around 200-300 ms post target appearance for distractor processing (Chelazzi et al, 2019; Sawaki & Luck, 2010; Falkenstein et al, 1999; Harper et al, 2014).

Whilst the time window of electrophysiological effects is similar to the distraction suppression literature, the pattern of results in this study differs greatly. In

contrast to the persistent effect of distraction during learning in this paradigm, many recent studies found improved suppression of probable distractor locations (Noonan et al, 2016; Gaspelin et al, 2015, 2017). However, the vast differences in task design make any direct comparison futile. To highlight some of the differences, previous studies have used a much greater number of trials (for example, above 2000 trials in Sawaki & Luck, 2010), and have not separated inhibition processes within visual search from memory-based negative bias, as this study attempted. The temporal separation of suppression acquisition and measurement is not an impediment in principle (see Turatto & Pascucci, 2016; Turatto et al, 2018), yet may complicate the evaluation of residual suppression effects if no stable suppression was established during learning.

There are a few considerations on why the Learning task may have failed to establish stable distractor suppression beyond the limited number of exposures. Firstly, the Learning task had equal numbers of scenes with and without distractors, which was required to contrast the effects of distractor association on memory precision and memory-guided attention. However, we know from the literature that even if distractor suppression can be acquired over time, attention capture can make a spontaneous recovery if the distractor suppression wasn't stable enough yet (Turatto & Pascucci, 2016). Due to the mixed design, no stable distractor suppression may have been able to be built up, especially noting the impact of block versus trial-by-trial design on the development of distractor suppression (Gaspelin & Luck, 2019). Secondly, the failure to suppress the distractor may be caused by the features of the distractor itself. I chose a highly salient distractor in terms of colour, and repeated and sudden onset (flickering), which may have been too attention grabbing to be suppressed (see Yantis & Jonides, 1984). The majority of recent work on distractor suppression utilised additional singleton paradigms, using a different singleton feature like colour as the "distracting"

dimension (e.g., Wang & Theeuwes, 2018a, b; Theeuwes & Van der Burg, 2011). It has been noted before that it is unclear whether the effects with such low-level feature distraction would translate to more salient distractors (cf. Gaspelin & Luck, 2019).

Beyond the constraints of the Learning task, there are further limitations of the study that have been alluded to in previous sections. A subsequent study would benefit from increased trial numbers, both for behavioural and electrophysiological data. However, and more importantly, there would need to be a greater number of learning blocks, extending the exposure to distractor-scene associations until the distractor can actually be suppressed. Achieving stable distractor suppression may require overtraining which might not be feasible in a classic in-person testing setup. Instead, the stable memory association could be achieved with a combination of testing strategies, using a remote setup for the Learning task and an in-person approach to the EEG measurement of the Memory-guided Attention task. In any case, the evaluation of negative memory biases and hence the possibility of an inhibitory tag intrinsically depend on the establishment of distractor suppression during learning. As previously discussed, the distractor features might have to be decreased in saliency to achieve suppression.

Furthermore, an interesting manipulation of the Learning task might include varying the location of the distractor. In the current design the location of the distractor was consistent within scenes, but changed on a trial-by-trial basis, which – as discussed above – may have prevented potential suppression from building up. It would be interesting to contrast the effects of a consistent versus variable distractor location during the Learning task. According to the Literature (see Noonan et al, 2016), we might expect better suppression for a consistent location than a variable one.

Additionally, a future alteration to the paradigm design overall might include an explicit recall for distractor memory. Such a distraction measure could give indication on the quality of the distractor localisation, as well as information on its putative role as attractor mechanism. The distractor recall would have to be placed after the present Memory task in order to avoid increasing the task-relevance of the distractor in the participant's mind.

In conclusion, I have observed that memories can be used as salient guides for attention in spite of distractor associations. Even though the distractor was consistently affecting search performance and could not be suppressed, I found evidence for subtle remaining distractor-associated effects with both Memory and Memory-guided Attention tasks. The indications of the formation of a distractor association could also be observed in the electrophysiological data, showing effects within a time frame consistent with other distractor suppression literature. However, it was not clear whether the association represents an inhibitory tag, as it also demonstrated some attraction features. The pattern of results might be more clear-cut in a design with more trials and stronger memory associations and actual suppression of the distractor within the Learning task.

Chapter III. Selective reactivation during consolidation

Introduction

Previous studies have shown how our long-term memories can guide our attention to relevant locations in space (Summerfield et al, 2006; Stokes et al, 2012; Rosen et al, 2014; Rosen et al, 2015a, 2015b). When using memories, rather than perceptual information, to orient attention the involvement of the hippocampus stands out (Summerfield et al 2006). Involvement of the hippocampus during memory-guided attention is not surprising given that the hippocampus has long been identified as an essential region for memory, with much research dedicated to characterising its functional role (Squire & Knowlton, 1995; Byrne, 2017). Emphasis was placed on a role for the hippocampus in declarative (Squire et al, 2004), explicit (Graf & Schacter, 1985), and episodic memory (Burgess et al, 2002), as well as spatial (O'Keefe & Nadel, 1978), relational (Cohen & Eichenbaum, 1993; Konkel & Cohen, 2009), and contextual memory (Chun & Phelps, 1999; Goldfarb et al, 2016). Further, all three main subprocesses of memory, namely encoding, consolidation, and retrieval, exhibit hippocampal involvement: Famously, medial temporal lobe resection including the hippocampus in patient H.M. as treatment for his intractable epilepsy left him with profound anterograde (and a degree of retrograde) amnesia (Scoville & Millner, 1957; Corkin, 2002). However, the hippocampus has not only been shown to be involved in the formation of memory, yet also during consolidation (McClelland et al, 1995; Born & Wilhelm, 2012), and retrieval processes (Hashimoto et al, 2012; Duncan et al, 2012; see Aly & Turk-Browne, 2017).

Hippocampus and hippocampus-dependent memories have been shown to interact with attention, with attention modulating hippocampal processing and

hippocampal memories influencing attention (for review, see Aly & Turk-Browne, 2017). Eye movements, a marker of attentional deployment, can be guided by hippocampal memories, even when the memories themselves are not conscious (see Hannula et al, 2010, for review). In an associative memory task, participants showed a gaze bias towards the memorised face stimuli, with increased hippocampal activity predicting the subsequent correct face fixation (Hannula & Ranganath, 2009). This effect persisted irrespective of whether the correct stimulus was ultimately chosen in an explicit recall task. Eye movements can therefore reflect implicit relational memory, known to be dependent on the hippocampus. The causal role of the hippocampus on these eye movements can be inferred from the finding that amnesic patients with medial temporal lobe damage affecting disproportionately the hippocampus did not show the gaze bias effect described above (Hannula et al, 2007). Indeed, the influence of hippocampus-based memories on attention was corroborated in a multitude of paradigms such as contextual cueing (Chun & Phelps, 1999), previewing scenes (Hollingworth, 2009), change detection tasks (Rosen et al, 2014), and memory-guided attention (Summerfield et al, 2006). Hence, in principle, any hippocampal-related mechanisms improving memory could also modulate and improve the use of hippocampal memories in attentional guidance.

The hippocampus-based memories discussed are long-term memories, thought to persevere over temporal delay. Some of the memory-guided attention tasks described above actually separated the learning of mnemonic episodes and the testing the resulting memory and attentional guidance by a night of sleep (Summerfield et al, 2006). Testing on consecutive days introduced not only a long-term delay for the memory, but also allowed for the potential of sleep-based consolidation of the memories. However, this was not manipulated systematically and was rather a by-

product of the experimental design. Could the memories used to guide attention be influenced by and consolidated during sleep? To my knowledge, no study has yet attempted successfully to trace how the consolidation of long-term memories might affect their use as guides to attention.

An essential consideration for this line of enquiry is whether sleep impacts the form of memory relevant for memory-guided attention. Sleep has long been known to improve various forms of memory, such as the retention of non-words (Jenkins & Dallenbach, 1924). Notably the beneficial impact of sleep is not uniform (see Stickgold, 2013), but instead particularly pronounced in types of memory with strong hippocampal dependence: Sleep has been shown to improve explicit memory as opposed to familiarity (Drosopoulos et al, 2005), context memory rather than item memory (Van der Helm et al, 2011), and relational memory (Ellenbogen et al, 2007). Sleep has also been found to strengthen memories of the temporal sequence in picture sequences (Griessenberger et al, 2012). The relevance of the hippocampus is further emphasised by the disruptive effects of cortisol, which has strong effects on the limbic hippocampal circuit (Wilhelm et al, 2011a). Hence, the location memories within scene context used in memory-guided attention tasks, which are contextual and episodic in nature, should fall within the spectrum of memories influenced by sleep-related consolidation.

The process by which memories are stabilised progressively once acquired and thus preserved is called consolidation (Dudai, 2004). Traditionally, consolidation was an answer to the stability-plasticity dilemma: The brain should be stable enough to preserve relevant information and plastic enough to allow novel information to shape cognition, without these two systems interfering with each other by overriding information ("catastrophic interference") or blocking new learning. The standard two-stage model of memory "solved" this dilemma by attributing the roles of fast, temporary

store to the hippocampus and stable, long-term store to the neocortex (Marr, 1971; McClelland et al, 1995). The encoded memory traces in the hippocampus would get repeatedly reactivated and the corresponding traces in the neocortex strengthened over time, a process termed system consolidation (e.g., Born & Wilhelm, 2012). System consolidation is dissociable from synaptic consolidation (Dudai, 2004): Whilst synaptic consolidation refers to the brain's mechanisms of storing information within minutes on a neuronal level for example by the strength of synaptic connections, systems consolidation denotes the specific process by which hippocampal-dependent memories are gradually reorganised, and transformed, even potentially becoming independent of the hippocampus. Systems consolidation thus undeniably relies on synaptic consolidation of the involved neuronal circuits. It is important to note that the transformation of memories from hippocampus to neocortex might lead to qualitative changes in the memories and does not necessarily entail a replacement of the original memory (Winocur et al, 2010).

There is widespread agreement that sleep can facilitate the consolidation of newly learnt memories (Maquet, 2001; Stickgold, 2005; Born et al, 2006). The role of sleep in memory consolidation is outlined in the influential Active Systems Consolidation hypothesis (e.g., Born & Wilhelm, 2012; Rasch & Born, 2013). It stipulates that sleep does not merely provide a protective period from retroactive interference (Ellenbogen et al, 2006), but that sleep physiology actively supports and promotes the consolidation process (Marshall & Born, 2007). The physiological processes involved in consolidation are described as follows (Inostroza & Born, 2013): Reactivations of mnemonic episodes originate within the hippocampus before spreading into the neocortex. The communication between the neural networks is controlled and enabled by oscillations: The depolarising up-phase of cortical slow oscillations in slow wave

sleep (N3) drives physiological events in the hippocampus and the thalamus. The hippocampus in turn reactivates memories and produces concurrent sharp wave ripples. Due to the common top-down drive from the slow oscillations, the hippocampal ripples occur simultaneously with spindles originating in the thalamus, forming spindle-ripple events. The ripples, carrying the reactivated memory information, become embedded in the excitatory down-phase of the spindle oscillation. These spindle-ripple events allow the reactivated memory information to be propagated to extrahippocampal, mainly neocortical areas. The process driving the spindle-ripple formation and propagation still occurs within the up-phase of the slow oscillation, which entails that the cortical neurons are in an excitable phase when the signal reaches them. The calcium-channel gated intracellular mechanisms enable synaptic consolidation and thus the long-term storage of the information within cortical areas.

Active System Consolidation thus proposes that sleep uniquely drives the reorganisation and strengthening of memory traces. Beyond the reactivation of memories during sleep, the Active System Consolidation hypothesis also makes certain, testable predictions about the consolidation process (Born & Wilhelm, 2012): (1) Memories reactivated during sleep are consolidated, (2) consolidated memories undergo qualitative changes, and (3) the consolidation process is selective. A landmark study on spontaneous memory reactivation during sleep showed that the patterns of spiking activity in hippocampal subregion CA1 in rats correlated while awake and searching for food and during subsequent slow wave sleep (Wilson & McNaughton, 1994). Importantly, the patterns of the activation during sleep followed the temporal sequence of firing during the spatial exploration (Skaggs & McNaughton, 1996) and were compressed in time, thus the sequence was replayed faster during sleep (Nádasdy et al, 1999). When hippocampal ripples were selectively suppressed during the

consolidation periods post training, the rats showed impairment in their spatial memory (Girardeau et al, 2009). Spontaneous replay during periods of rest and subsequent effects on memory were also observed in humans (Zhang et al, 2018; Liu et al, 2019). Multiple studies evidenced qualitative changes in the memories, such as the disappearance of idiosyncratic details (Nadel & Moscovitch, 1997), and greater assimilation and abstraction (Lewis & Durrant, 2011; Ellenbogen et al, 2007). There is also evidence that the consolidation of memories is selective, preferring memories with future relevance (Wilhelm et al, 2011b) or associated reward (Fischer & Born, 2009).

How can we utilise the observations that reactivated memories are consolidated, and that consolidation is selective, in order to influence which memories are consolidated preferentially? A methodological solution to this question was presented in form of targeted memory reactivations (reviewed by Oudiette & Paller, 2013). Rasch and colleagues (2007) were first able to use odours to cue reactivation of object-location associations in human participants during slow-wave sleep. Cueing the odour context improved the retention of the associated hippocampal-based memories, in the absence of a comparable effect on hippocampus-independent procedural memories. Rudoy and colleagues (2009) expanded this finding to cueing individual target-location association with auditory cues, resulting in higher memory accuracy for cued associations in comparison to uncued ones. Targeted memory reactivation is thought to help memory consolidation by reactivating memories using external cues (Oudiette & Paller, 2013). External cues are often used in recall tasks and are known to trigger a process of pattern completion (Rolls, 2013) in the hippocampus, by which the target memory can be reinstated in the neocortex (Staresina & Wimber, 2019). How a cue is able to bias hippocampal reactivation in sleep is yet to be described, as are the optimal conditions for applying targeted memory reactivation (Schouten et al, 2017). Regardless, there is

evidence recorded from hippocampal cells in rats, that auditory cues during sleep biased reactivation towards replaying the relevant events of the trained auditory-spatial associations tasks (Bendor & Wilson, 2012). The efficacy of targeted memory reactivation, under certain circumstances, has since been replicated widely (for example, Antony et al, 2012; Oudiette et al, 2013; Creery et al, 2015; Vargas et al, 2019; Schechtman et al, 2020). A recent meta-analysis on targeted memory reactivation found strong evidence that it can influence memory storage when applied during non-REM sleep Stages 2 and 3, and supported the method as a means to investigate memory consolidation (Hu et al, 2020).

Circling back to the question about the impact of memory consolidation on memory-guided attention, multiple conclusions about potential methodological approaches can be drawn from the evidence above. It has shown that the hippocampus is involved in memory encoding, consolidation, and retrieval; and that hippocampal-based memories interact with attention. Indeed, the location-scene associations commonly used in memory-guided attention tasks show hippocampal involvement. Hippocampal-based memories were shown to be consolidated during sleep. This consolidation appeared to be causally dependent on hippocampal reactivations, which in turn can be triggered by external cues, for example in a targeted memory reactivation paradigm.

Accordingly, I adjusted the memory-guided attention paradigm (Summerfield et al, 2006) to pair auditory tunes with the learning of target locations within scenes during the Learning phase. Half of these idiosyncratic tunes were replayed during a nap phase when participants were judged to be in slow-wave sleep, aiming to trigger preferential reactivation and thus consolidation of the associated memory pairings. A Memory task in form of a direct recall task was included before and after the nap phase to evaluate

the presence and quality of memories and to measure any changes in memory resulting from to sleep and auditory stimulation during sleep. The Memory-guided Attention task was designed to measure how memory cues can be used to guide our selective attention within space (Summerfield et al, 2006). Within the current adapted version of the task, I was able to contrast the effects of sleep, and sleep-based stimulation on attentional guidance. Will the prevalence of slow-wave sleep and targeted memory reactivation influence the accuracy with which we store memories and the speed with which we access them? And will they affect the way we use our memories to guide attention?

The nature of the task offered itself to auditory tunes (as opposed to olfactory stimuli) for sleep-based stimulation. However, in contrast to many targeted memory reactivation studies using auditory cues (Rudoy et al, 2009; Oudiette et al, 2013), in this study the auditory stimuli were not matched to the visual targets on a one-by-one basis. Instead, each tune was associated with a scene type, or context, such as beach or forest. The context pairing in this study was reminiscent of the olfactory context chosen in the original targeted memory reactivation study by Rasch and colleagues (2007): Can recreating the auditory context during sleep lead to similar preferential memory consolidation? The large number of scenes, and consequently of required tunes, also worked against the possibility of individual pairing: Due to the few learning exposures, it was questionable whether participants would have developed distinct memories for each of the tunes.

This study allowed us to focus on the following three strands of enquiries, beyond replicating the standard memory-based attention effects. Firstly, I wanted to determine whether the memory utilised in memory-guided attention was subject to sleep-dependent consolidation. I hypothesised that sleep decreased forgetting, and that it emphasised the validity effect observed in the memory-guided attention task. This

would occur if RTs were decreased for memorised, valid locations, and/or increased for invalid locations for longer sleep durations. Secondly, I was interested to test whether auditory stimulation in sleep could bias memory consolidation even if the sounds are not conceptually linked to the reactivated memory, but abstract and merely associated during encoding. Hereby it is important to note that these tunes were more complex than usual and were associated to categories of scenes. I hypothesised that there would be less forgetting for stimulated scene categories. Thirdly, if there was an effect of stimulation on sleep-dependent consolidation and subsequent memory performance, I wondered whether this would also affect the use of the memory to guide attention. I predicted that sleep-dependent stimulation would increase the validity effect.

Method

Participants

Twenty-nine healthy and right-handed volunteers participated in the study. All participants had normal or corrected vision and reported no difficulties with colour vision. Further, participants reported no history of neurological or psychiatric disorders and were on no medication that might affect their cognition. In total, data of six participants had to be excluded. Four participants were excluded because the experiment was not completed (2), or difficulties with a task (1), or an error with the sleep manipulation (1). One further participant was excluded due to having a sleep disorder. Another participant was excluded due to poor memory performance.

Consequently, the data of 23 participants were analysed (age range: 18 - 30; mean age = 23.67; *SD* = 3.37; 15 female; age and handedness from two participants are

missing). The study was approved by the University of Oxford Central University Research Ethics Committee (CUREC) and conducted in accordance with the Declaration of Helsinki. All participants gave written informed consent before testing and were financially reimbursed for their time (£65).

Materials

Visual stimuli

One hundred and forty-four naturalistic scenes were used in the experiment. These scenes were photographs from four different natural categories: Forests, house fronts, indoor displays, and beach scenes. All images had the dimension of 1920 x 1080 pixels. With a viewing distance of 100 cm, the visual angle of the stimuli was approximately 33ºx19º. For the Learning task only, a small gold-coloured key (55 × 25 pixels; 0.98º × 0.45º visual angles) was placed in one of the four quadrants. In order to counterbalance target locations across experimental conditions in the study cohort, for each scene, the target could be located in one of four locations defined according to quadrants (by dividing the screen halfway both horizontally and vertically). Keys were preferably placed in a hidden location within the designated quadrant to encourage search and memory formation. A further, larger, key image (60 x 28 pixels, 1.07º × 0.50º visual angles), was used as target in the Memory-guided Attention task.

Auditory stimuli

Further, the study utilised auditory stimuli, an assortment of tunes and background noise. All auditory stimuli were presented via sound isolating headphones (Sennheiser HD449).

Brief auditory tunes to pair with learning of different scene categories were specifically composed for the study using Sibelius 7 (sampling rate: 44.1 kHz, 16 bits per sample). These tunes were composed in collaboration with the musicologist and composer Dr. Lewis Coenen-Rowe. They were specifically designed to be distinctive in sound and loop smoothly to create a homogenous auditory background theme. The distinctiveness was important to make them recognisably different even for participants with an unschooled musical ear, as piloted by lab members. To achieve such individuality, the tunes varied in sound quality, synthesised instrument use, pattern and speed (*Figure 10*). Tune 1 had the duration of 10s, 133 BPM, and was played by strings and clarinet in 3/4 meter. Tune 2 had the duration of 13s, 64 BPM, and used strings and choir in 6/8 meter. Tune 3 had the duration of 13s, 121 BPM, and used an oboe, cor anglais, and bassoon in a 2/4 meter. The duration of tune 4 was 32s, with 120 BPM, and the tune was played by two violins and a violoncello in a 6/8 meter. The tunes were presented unchanged during the Learning task, yet attenuated to 15% of the maximum volume control at the computer level.
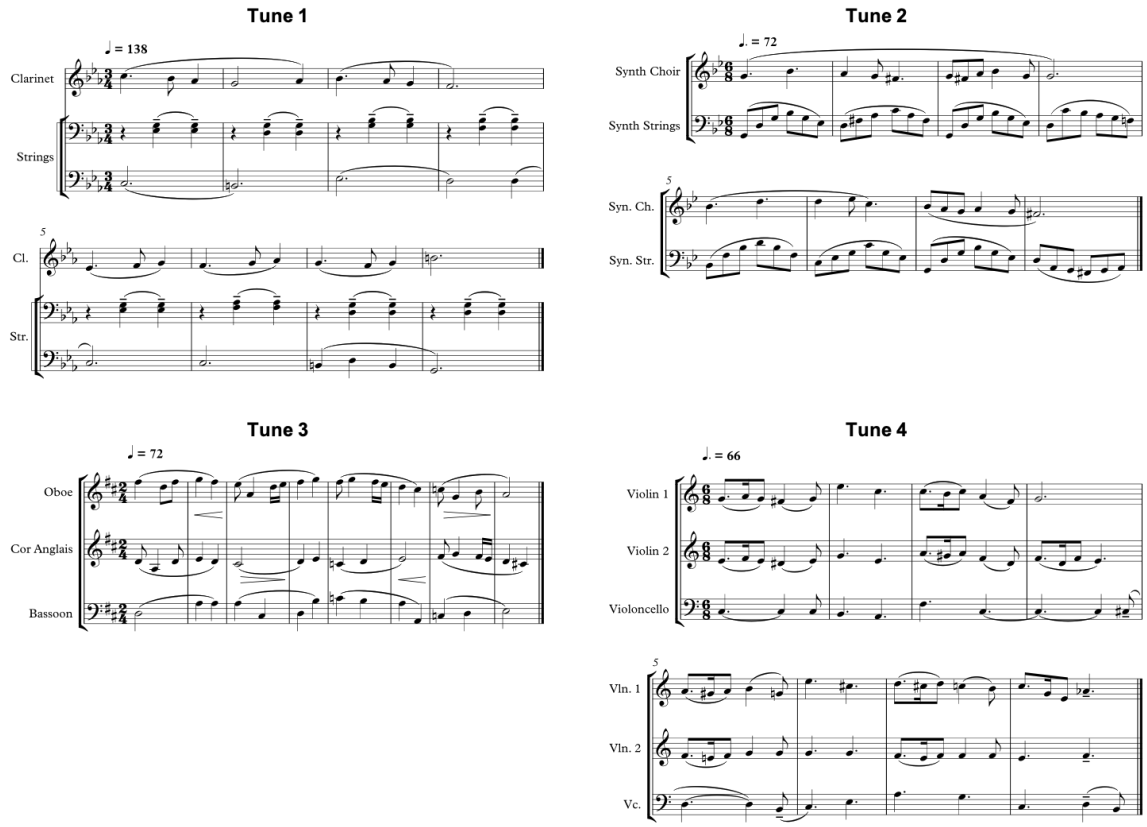


Figure 10. The sheet music for the four different tunes is presented. Each tune differs in speed, sound quality and instruments used.

The second set of auditory stimuli, required for the Nap phase, were altered versions of the first set. The four tunes were each cut to the length of 10 s, and a two-second fade in and a two-second fade out was implemented for each of the tunes. The fade-in was required to decrease arousal risk in response to a sudden change in volume. The resulting BPMs for the tunes were as follows: Tune 1 111 BPM, Tune 2 117 BPM, Tune 3 117 BPM, and Tune 4 120 BPM. The loudness level of the auditory stimulation was required to remain under the arousal threshold of around 30 decibels. The threshold was estimated by the mean awakening and arousal threshold for N3 minus two times the standard deviation in normal sleepers (Pilon et al, 2012). The loudness level of each tune was hence attenuated using the presentation software (Tune 1: 0.637, Tune 2: 0.603, Tune 3: 0.638, Tune 4: 0.627). Decibels were measured using a decibel meter

(TSSL1 sound level meter). In addition, a brown noise stimulus was used for the Nap phase (sampling rate 48 kHz, 16 bits per sample). The brown noise was normalised to smooth out potential random differences in noise levels and cut to the length of 10 minutes. Attenuation was set using presentation software at 0.512. The auditory stimuli used in the Nap phase were presented at a computer volume setting of 50% and could be presented simultaneously (tune and brown noise). Each of these stages, namely the reported arousal threshold, the attenuated tunes, and the attenuated tunes in combination with the brown noise stimulus, was tested in pilot participants to evaluate whether the sound level caused arousal.

One participant struggled with the sound level during the Learning task, finding the volume too loud, so the computer volume control was decreased from 15% to 10%. Equivalently, the volume of the (already digitally attenuated) auditory stimuli during the Nap phase was further attenuated for this participant from 50% to 45% using the computer.

Sleep set-up

For the Nap phase participants were asked to lie on a bed or mattress in the shielded booth, equivalent to a different room. Three different sleep environments were used across the cohort, but all had equivalent properties. Participants were given a blanket and pillows for comfort. One pillow was placed underneath the knees to pre-empt participants turning or moving too much in their sleep. A towel was placed on the head pillow to catch any spillage of the EEG gel. Participants were also given headphones to wear, both as a noise-shield to reduce the risk of arousal due to ambient noise, and as a means of presenting the auditory stimulation covertly during sleep.

Experimental set-up

During the experimental tasks, eye-gaze position was measured using a remote, infrared eye-tracking system (Eyelink 1000). To facilitate eye tracking and stabilise head position for the tasks, a chin rest was used throughout the experiment. The height of this chin rest was adjusted for each participant. This set-up was held constant across experimental tasks. EEG was recorded during the first Memory task and the Nap period.

The experimental tasks and the stimulus presentation during the Nap period were controlled via the Presentation software (version 18.3; Neurobehavioral Systems). Presentation was also used for recording behavioural responses and synchronising experimental timing with the eyetracking and EEG systems. All tasks were presented on an Acer XB270H Monitor (27-inch, resolution: 1920 x 1080 pixels; screen width: 60 cm), using a refresh rate of 60 Hz. The statistical analyses were run with SPSS (Version 26.0.0.0) or R. For each statistical test, the respective assumptions were checked and a correction for multiple comparison (Bonferroni) was applied where appropriate.

Design

In the experiment, participants performed three distinct tasks: Learning, Memory, and Memory-guided Attention. The Memory task was performed twice in its entirety. The tasks are illustrated in *Figure 12-14* a.

Learning task

The aim of the Learning task was to establish new long-term memories for spatial locations within naturalistic scenes, *Figure 12a*. Participants viewed 144 scenes containing a small hidden target in shape of a key in a randomized order.

The instructions were to search through each of the scenes to find the hidden target key. One second prior to scene onset, one of the four tunes started. Each tune was associated with one scene type for each participant, but this was not explicitly stated to the participant. After scene onset, participants were free to move their eyes, and their gaze position was recorded. Maximum search time was set at 90 s. Once participants had found the key, participants had to click once on the mouse to activate a cursor. This cursor could then be moved freely around the scene to enable participants to click on the location of the found key. The maximal time to move the cursor to the key location was set at 5 s. Time until the first click was counted as reaction time, so participants were instructed to click only after they had spotted the key. The tune disappeared abruptly after the first mouse click. Participants were given feedback on screen (500 ms). If their click was within 50 pixels of the actual target location (measured in Euclidean distance), participants were given the positive feedback "Target found". When participants clicked too far off the target or did not respond within the maximum search or response window, the feedback screen read "Target Not Found". The inter-trial interval (ITI) was randomly drawn between 1000 and 1500 ms. Further feedback was provided at the end of each of the five learning blocks, giving the participants information on how many keys they had found in this particular learning block. For any given participant, the specific location for the target within a particular scene remained constant across learning blocks as did the association between each tune and a scene category. However, both of these factors were counterbalanced across participants, to pre-empt spurious effects related to salience of target locations within individual scenes or scene categories (*Figure 12b*). The task design allowed for a within-subject analysis of effects of learning block (5 levels), scene category (4 levels), and tune (4 levels). Dependent measures included identification accuracy and reaction time.

Memory task

The Memory task was an explicit spatial recall task, *Figure 13a*. It was used to determine the presence and accuracy of memories for the target locations for scenes in different experimental conditions. The Memory task was completed once pre-nap and once post-nap. Participants viewed the 144 learnt scenes in a randomized order. No target or auditory stimuli were present for this viewing. Participants had to indicate the location of the remembered target associated with each specific scene.

Each scene appeared for a maximum duration of 20 s. The instructions were to click a mouse button once participants remembered the target location. Participants were then asked to move the activated cursor to the remembered location and indicate it with a second click (max 20 s). Participants were asked to give their best guess of the target location should they not remember. Lastly, participants had to rate their confidence in their memory using a three-point scale (1 = not at all confident, 2 = fairly confident, 3 = very confident). The confidence scale was presented on screen until a response was made. Trials were separated by an ITI of 900 ms with a central fixation square. As in the other tasks, participants kept a stable position using a chin rest, but they were allowed to move their eyes freely to explore the scene. Eye gaze was tracked throughout the first scene presentation. The main recorded dependent measures were RT and spatial error in Euclidean distance. During the pre-nap Memory task, EEG was recorded, with triggers sent for the scene onset.

Nap phase

During the nap period, participants were lying in the booth or adjacent nap room. Using the headphones, brown noise was played continuously to attenuate ambient noise. The experimenter monitored the EEG signal for abnormalities, arousal signs, and

sleep markers. Once participants were judged to be in deep sleep (N3 if possible, otherwise N2), the experimenter would activate the auditory stimulation protocol. Two of the tunes presented during the Learning task were replayed, the choice of which was balanced across participants. The numerically lower tune number was always presented first. After a delay of 24 s, the second tune was presented. Both replayed tunes were always presented in combination, with one presentation of each tune constituting one round of stimulation. After a further delay of 24 s, the first tune was presented again. The loop of delays and tune presentations was continued until the experimenter detected signals of arousal, sudden ambient noise, or other conditions that might lead to an arousal. The auditory stimulation could be triggered on multiple occasions during the Nap phase, and each stimulus presentation was marked with a trigger. If no stable sleep was detected, no auditory stimulation was activated.

Memory-guided Attention task

The Memory-guided Attention task was a Posner-like spatial orienting task (Posner, 1980) using the learnt scenes as spatial memory cues for the target location, *Figure 14a*.

Each trial consisted of the following elements: A previously encountered scene appeared briefly on the screen without an embedded key (100 ms). A blank 1200-ms interstimulus interval followed (ISI, grey screen with fixation square), after which the target key flashed up (100 ms) on a neutral, grey background. The key could either appear in the learnt position (condition: valid), or a different, unlearnt position (condition: invalid). The task was to detect the appearance of the target by clicking on the mouse as quickly as possible. The key did not stay on the screen and participants did not have to click on the target location. A blank inter-trial interval then began (ITI,

grey screen with fixation square, 1000 – 1500 ms), which doubled as response window, irrespective of response given. Participants were required to maintain central fixation throughout the Memory-guided Attention task, and they were given breaks at regular intervals (every sixteen trials) to rest their eyes.

Participants completed 112 trials, each containing one of the learnt scenes. (Thirty-two of the 144 learnt scenes were not presented in the Memory-guided Attention task in order to obtain a 'clean' measure of post-nap memory, uncontaminated by targets appearing at valid or invalid locations during the Memory-guided Attention task.) Sixteen of the 112 presented trials were catch trials, in which no target was presented on screen. These catch trials were used to encourage participants to pay attention to whether a target actually appeared, rather than clicking blindly after scene presentation. This was especially important due to the fixed ISI. Catch trials were excluded from the main analysis but were evaluated separately to get a measure of response bias to non-existent targets.

The remaining 96 trials were target trials: 48 trials for the valid condition and 48 for the invalid condition. Within valid trials, the key appeared at the same location it had been encountered within the relevant scene during the Learning task. In invalid trials, the key could appear in one of the other three quadrants. Overall, each quadrant was equally likely to contain a target. Half of each valid and invalid trials were associated with tunes replayed during the Nap phase, half were not. This led to 24 trials in each of four conditions: valid replayed, valid not replayed, invalid replayed, and invalid not replayed.

Procedure

Participants performed four experimental tasks and a nap period over two sessions within the same day. Prior to the testing day, participants were sent information about the study, a questionnaire on their usual sleep patterns (Morningness-Eveningness Questionnaire, Horne & Östberg, 1976), and some requests. The requests included trying to get enough sleep the nights preceding the study, and to have a fairly regular cycle of getting up and going to bed. This was meant to increase the likelihood that participants had their circadian dip around the time of the nap opportunity and increase chances of sleep during the nap opportunity. Further, participants were asked to refrain from caffeine (such as coffee, tea, or fizzy drinks) as well as stimulants (such as alcohol) on the testing day. Again, this was aimed to increase the chances of sleep during the Nap phase and decrease the likelihood of abnormal sleep.

On the testing day, participants were informed again that the study included four experimental tasks, over two sessions. Participants were informed of the use of eyetracking for some of the experimental tasks and the use of EEG during the first experimental session and the Nap phase. Participants had the chance to read through the study information once again, and then provided informed consent. Further, participants were given verbal and written instructions on the first task, the Learning task, including schematics. They were asked to find as many keys as possible and to memorize the exact location within the scene. The experimenter showed participants into the experimental booth, adjusted the chin rest, and calibrated eyetracking. After each block of the Learning task, the experimenter recalibrated the eye tracker and ensured that the participant was doing well. Participants performed five blocks of the Learning task, taking around two to three hours. Due to the nature of the study, every

participant was tested at the same time of the day. As the first experimental session started at 09:30 am, participants usually finished the Learning task around noon.

Subsequently, participants were given a bathroom and a lunch break. Lunch was either brought by the participant, or bought during the break, or provided for by the experimenter, depending on the request of the participant. Participants were again asked to refrain from stimulants that might affect their sleep. Depending on the time, EEG system was set up during the lunch break. The specifics about the EEG setup can be found in the section *EEG data acquisition*. Initially, participants' skin was prepared with an alcohol patch and an abrasive gel. Mastoid references and EOGs were placed before the remaining electrodes within the electrode cap. After the EEG set up, the experimenter explained the Memory task, using verbal and written instructions including schematics. The experimenter also explained the effects of certain behaviours on the EEG traces, such as grinding teeth, clenching of the jaws, blinking and moving eyes. Participants were encouraged not to do those motions during the EEG recording. When the set up was completed, participants moved back into the testing booth. The experimenter calibrated the eye tracker and participants then completed the first block of the Memory task (pre-nap), taking around 20 minutes. The benefits of setting up the EEG prior to the Memory task were to 1) provide delay between the Learning and the Memory task, and 2) give the electrode a chance to settle, and for unstable EEG traces to become visible before the Nap phase.

Once the first Memory task was completed, electrodes were placed on the mandibles and chin to measure muscle activity (EMG, electromyograph). They could not be placed before the Memory task due to the use of a chin rest. Participants were then helped onto the bed or bedlike structure. Their head was bedded on a pillow to fixate the head in position, and their knees on another pillow for a similar purpose.

Participants were informed that they would be wearing headphones for the duration of the Nap phase, which would be playing brown noise. The aim of the brown noise and the headphones was explained as masking ambient noise and reducing the risk of arousal. The secondary, main aim of the headphones as means of auditory stimulation was not mentioned. After the lights were switched off and the door closed, the brown noise and the EEG recording started.

The experimenter monitored the Nap phase by following the EEG, EOG and EMG traces. Once participants were judged to be in stable deep sleep (ideally N3) using online sleep scoring, the auditory stimulation was started. The experimenter paid close attention to any signs of arousal as well as ambient factors that might disturb the deep sleep. If there was a sign of arousal, or a perceived risk of potential arousal, the experimenter would disrupt the loop of auditory stimulation. The duration of the Nap phase was 90 minutes. Ideally, the beginning of the Nap phase was timed to start around 1:30 pm, varying slightly depending on the progress of the experiment. Around 3:00 pm participants were woken up by the experimenter or awoke naturally. Participants were then given time to wake up, and to wash, condition, and dry their hair. A two-hour break following the Nap phase aimed at decreasing any chance of sleep inertia.

At 5:00 pm the second experimental part of the study began. It included the Memory-guided Attention task and the post-nap Memory task. For half of the participants, the Memory-guided Attention task was presented first and for the remaining half the post-nap Memory task was presented first. For the Memory task, eyetracking was calibrated again. Prior to the Memory-guided Attention task, participants were given verbal and written instruction with schematics, as well as a practice task. The practice task was an equivalent example version of the Memory-guided Attention task using novel scene stimuli. It consisted of twelve trials including

four catch trials. The training task was implemented to help the participant learn to cope with the physical requirements of the task (e.g., maintain central fixation and not to blink) as well as to ensure they only clicked after they detect the transient key target, avoiding being caught out by catch trials.

Behavioural Analysis

Learning task

Learning was measured by the accuracy and RT of target identification over the course of learning blocks. For statistical analyses, I transformed the data using arcsine for accuracy and log for RT data. Trial exclusion was based on accuracy in detecting the key and RT outliers beyond three standard deviations of the mean of the respective block and condition. I analysed data for all participants with regards to effects of block (5 levels) using Friedman test or repeated-measures ANOVAs as appropriate. Progressive learning was assessed with a linear contrast of the block factor. Separately, I analysed effects for block on accuracy and RT for a subset of participants, defined by having received at least one round of stimulation during slow-wave sleep (n=15).[2]

Memory task

Memory performance was assessed using RT and spatial error in Euclidean distance, as well as confidence ratings. Log data transformations were used for RT and spatial error data. Only trials without response were excluded from the analysis. Spatial

[2] For this study, all analyses were intended to be run on 24 participants, each of which having spent enough time in slow-wave sleep to receive auditory stimulation. Due to the current situation and restrictions on in-person testing, I was not able to ensure that enough participants meeting these requirements were tested.

error data were compared to chance level, estimated using a permutation approach outlined in the previous chapter (Chapter II, Section: Methods. Behavioural Analysis. Memory task.). Memory performance based on spatial error and RT data were contrasted over the two sessions, pre- and post-nap, to assess effects of forgetting. These analyses used trials unbiased by sleep-based stimulation (non-replay). Further, spatial error and RT data from a subset of participants (n=15), having received a minimum of one round of stimulation in slow-wave sleep, were analysed based on Memory session and according to whether the tune associated with the scene category was replayed during sleep. These conditions resulted in two 2 (session) by 2 (replay) repeated-measures ANOVAs. In an independent analysis, I separated RT and spatial error data with respect to associated confidence scale. This was assessed by two repeated-measures ANOVAs using session (two levels) and confidence score (three levels) as factors.

Further, I correlated physiological markers of sleep with performance in the memory task: For this, I firstly used the spatial error post-nap in scenes not associated with replay across all participants. Spearman's rank correlations were used on transformed data to minimise the effects of potential outliers. As outlined below, these sleep markers consisted of the overall duration in sleep and the duration of individual sleep stages. Secondly, I analysed correlation of sleep data with a normalised forgetting score, calculated by the difference of mean spatial error pre- and post-nap, divided by the sum, for each individual subject: $\frac{M_{post-nap}-M_{pre-nap}}{M_{post-nap}+M_{pre-nap}}$. For the sleep-dependent correlations I used the normalised forgetting score from non-replayed trials only.

Lastly, I assessed the effects of replay on memory performance: In Spearman's rank correlations, I correlated the spatial error data post-nap (of scenes associated with

replay) with the duration of stimulation (in N3). I used transformed data and performed the correlations over all participants with a minimum of one round of stimulation (n=15). Effects of stimulation duration were also correlated with the normalised forgetting score in replayed scenes, as well as the differential forgetting effect (normalised forgetting in replayed trials – normalised forgetting in non-replayed trials). I also ran exploratory analyses on a subset who received more stimulation, namely at least five rounds of stimulation (n=12), given that very few rounds of stimulation might not be sufficient to induce targeted memory reactivation.

Memory-guided Attention task

The analysis of the Memory-guided Attention task focused primarily on RT data. Accuracy data were merely reported to illustrate compliance with and viability of the task. For statistical analyses, RT data were log transformed. Exclusion of data was based on accuracy (response given after target appearance) and RT outliers beyond three standard deviations of the condition mean for each participant. Firstly, RT data were analysed solely with respect to cue condition over all participants, using only trials not associated with replay. Secondly, RT data of the subgroup of participants with minimum of one round of stimulation (n=15) were analysed with respect to cue condition and replay: Targets could appear in a learnt (valid) or unlearnt (invalid), and could be associated with sleep-based stimulation (replay) or not (no replay), resulting in a 2-by-2 repeated-measures ANOVA.

For all following correlations, Spearman's rank correlations were chosen due to the data distribution and for their robustness against outliers. A normalised validity score was calculated to allow correlations using a single value per participant: $\frac{M_{invalid}-M_{valid}}{M_{valid}+M_{invalid}}$. This validity score was correlated with the post-nap memory

performance (spatial error in Memory task) to assess the relationship of memory quality on its use for attentional guidance. Further, the validity score was correlated with measures of sleep physiology and stimulation: Equivalent to the Memory analysis outlined above, correlations focused on the sleep measures used validity values from scenes not associated with replay. In contrast, correlations with the duration of auditory cue stimulation used validity values from replayed trials. Additionally, I correlated the differential cueing effect (validity effect of replay – no replay) with the duration of sleep-based stimulation, assessing whether stimulation conferred a benefit on the use of memories on guiding attention. I also completed the exploratory analyses on a further subset (minimum of five rounds of stimulation, n=12) with higher stimulation intensity for the last set of analyses.

EEG Data Acquisition

EEG data were collected with Synamps amplifiers and data acquisition software by Neuroscan (Compumedics Neuroscan, El Paso, TX) from 48 scalp sites, using Ag/AgCl electrodes mounted on an elastic cap, positioned according to the 10-20 international system (AEEGS, 1991). Mastoid references were used, the active reference placed on the left mastoid bone. Data were later re-referenced off-line to the algebraic average of the two mastoid reference electrodes for the vast majority of participants. In one exception, data were re-referenced to Cz due to the disconnection of a mastoid electrode during the nap. The common ground electrode was placed in the cap in the position of FpZ. The montage consisted of 8 midline electrode sites (AFz, FZ, FCZ, CZ, CPZ, PZ, POZ, and OZ) and 20 sites over each hemisphere (AF3/AF4, F5/F6, F3/F4, F1/F2, FC5/FC6, FC3/FC4, FC1/FC2, C5/C6, C3/C4, C1/C2, CP5/CP6, CP3/CP4, CP1/CP2, P7/8, P5/P6, P3/P4, P1/P2, PO7/PO8, PO3/PO4, O1/O2). In

addition, vertical and horizontal Electrooculograms (EOG) were recorded. The vertical EOG was used for detecting blinks and was derived by subtracting signals between two electrodes placed above and below the right eye. The horizontal EOG was used to detect saccades and was derived by subtracting signals from electrodes on the outer side of right and left eyes. Electromyography (EMG) was recorded with one electrode placed on each side of the sub mentalis muscle below the chin, and a third, reference, electrode placed on the chin. All EEG impedances were kept ideally below 5 kΩ (at least below 10 kΩ), EOG and EMG below 10 kΩ. Signals were recorded with an acquisition rate of 1000 Hz, a low-pass filter of 200 Hz and a high-pass filter of 0.1 Hz. Raw data of two participants had to be merged digitally, as recording was interrupted during the Nap phase and needed to be restarted.

EEG Analysis/Sleep scoring

The quality of the raw EEG data was assessed via power spectra and visual observation. Bad, or disconnected channels were interpolated using the spherical method. EMGs, EOGs or reference channels were not interpolated but rather a replacement found in terms of sleep scoring. All data were re-referenced off-line to the algebraic average of the mastoid references (or Cz for one participant, for whom it was not possible to derive a linked mastoid reference, after confirming the reference method left sleep markers intact). All data were transformed using SPM12 (Wellcome Centre for Human Neuroimaging, London) into .dat format. Sleep scoring was performed using the FASST toolbox (Schrouff et al, 2009; Leclercq et al, 2011). The data were scored visually by three trained scorers, following the criteria set by the AASM Manual for the Scoring of Sleep and Associated Events (Iber et al, 2007).

Identifying and scoring sleep stages

The sleep-scoring criteria followed those outlined according by Malhotra and Avidan (2013). *Figure 11* shows some representative samples in which the distinctive features guiding scoring are labelled. An epoch was scored as awake if more than 50% of the epoch consisted of scorable alpha (label 2), included high muscle tone or a major muscle event (label 4) or arousal. Slow rolling eye movement (label 6), slowing of the alpha frequency or emerging theta rhythms (label 5) are signals of participants falling asleep, with the sleep score of NREM Stage 1 (N1) for light sleep. During light sleep, participants are easily aroused, for example by noises or even involuntary muscle twitches (hypnic jerks), and the transitions are often fluid. This makes the transitional stage of N1 difficult to score. Especially within this nap setting, participants might drift off and awake often. We thus determined the time of "sleep onset" as representing the first instance of N1, where the next stage was deeper sleep (NREM Stage 2 or N2) rather than wakefulness again. N2 still contains theta rhythms, the predominant oscillatory frequency in sleep. Markers of sleep that set N2 apart are k complexes (label 7) and sleep spindles (label 8). K complexes are sharp slow waves of a duration of at least 500 ms, showing a negative followed by a positive inflection. They are the largest signal the healthy brain produces, and can occur spontaneously or evoked, for example by noises (Cash et al, 2009). Sleep spindles are prolonged oscillations in the sigma band (11 – 16 Hz), lasting at least 500 ms. They are notable for a high degree of symmetry and synchrony, creating a "spindle"-like appearance in the EEG trace. Spindles have been associated with memory consolidation (Diekelman & Born, 2010). According to functionality and topography, a distinction of frontal and central spindles has been proposed (e.g., Vargas et al, 2019). Sleep stage N2 is scored if such a k complex or spindle occur within the first half of the epoch, or within the second half of the

preceding epoch. N2 is continued to be scored regardless of the presence of further spindles or k complexes, unless arousal occurs or other markers indicating transition to deeper sleep stages. N2 accounts for the majority for sleep.

The deepest Non-REM sleep stage is slow wave sleep, or N3. As the name implies, the marker of N3 are slow waves (9) within the delta band (0.5 – 2 Hz). Slow waves are predominantly seen in the frontal region and are required to have an amplitude of at least 75 μV. N3 is scored if at least 20% of the epoch demonstrates slow waves. Slow waves, and N3, have been associated with driving memory consolidation, which makes this sleep stage particularly interesting for the cognitive studies of memory.

Rapid-Eye Movement, or REM, sleep is the deepest sleep stage. It occurs about 90 – 120 minutes after sleep onset and is more pronounced in the sleep cycles of the second half of a night's sleep. It is also termed paradoxical sleep due to its EEG physiological profile: Despite being deeply asleep, participants show an EEG profile similar to wakefulness. The brain waves are small and irregular, showing theta and slower alpha. Characteristic are the sudden bursts of eye movements (10), and a very low muscle tone (11), stemming from muscle paralysis during this stage (atonia). REM sleep is often associated with vivid dreaming. REM stage is scored when the EEG signal decreases in amplitude, usually from N2 or N3, with a sudden drop in muscle tone and rapid deflections in the eye channels. Sleep scorers conventionally continue scoring REM, as was done here, irrespective of the presence of rapid eye movements, until arousal or a change in EEG traces becomes obvious.

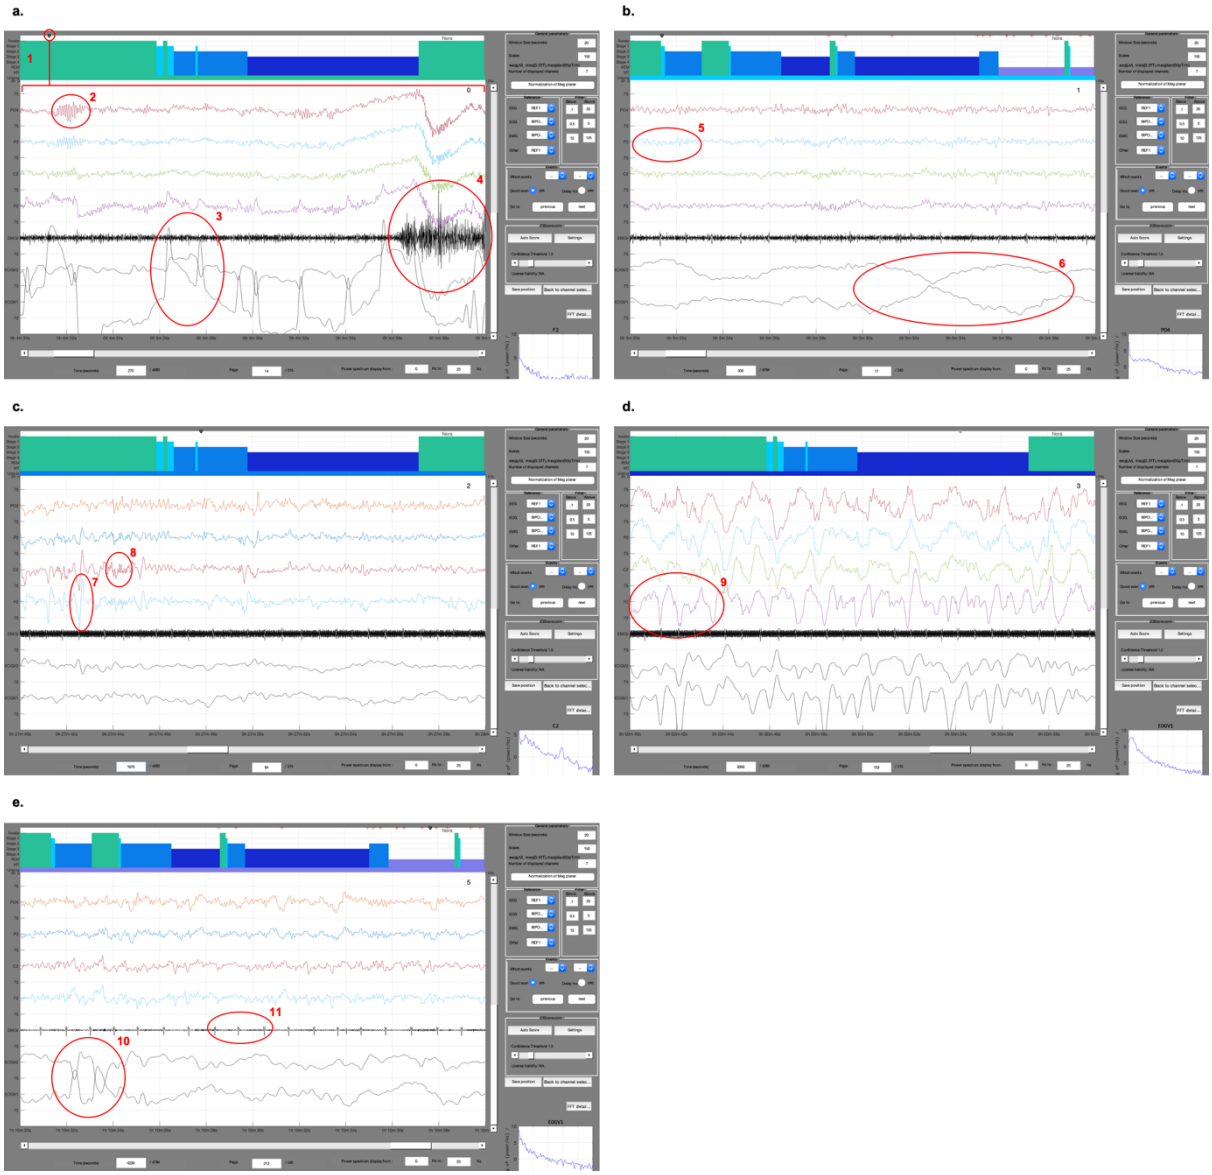


Figure 11. All panels are of the FASST scoring tool. Marker 1 indicates the blue and green hypnogram that develops for each participant and scorer. Here it is visible that panels a., c., and d. are from the same participant and scorer, and b. and e. are from a second participant. The marker at the top of (label 1) denotes when during the recording the EEG scores are taken from. Below the hypnogram are the EEG, EMG and EOG traces for selected electrodes. On the right side, the scoring window, amplitude scales and reference can be altered. In the bottom right, a power spectrum gives indication about predominant frequencies. Panel a. shows a participant during wakefulness, demonstrating alpha rhythm in the posterior channels (2), pronounced and rapid eye-movements likes saccades or blinks (3) as well as motion (4). Panel b. illustrates a typical profile for NREM stage 1: There are strong theta rhythms (5) and rolling eye-movements (6) as indicators of transitioning into sleep. Panel c. is taken from sleep stage NREM 2, indicated by k complexes (7) and spindles (8), both in combination and separately. Panel d. exemplifies a sleep scoring window for NREM Stage 3 or slow wave sleep, as slow waves (9) are broadly visible across all EEG channels, showing delta frequencies with amplitudes larger than 75 µV. Note: Sleep stages in this study were scored according to the recent guidelines, which combined stages N3 and N4 into one stage (N3). The FASST toolbox still includes a differentiation between these two stages. We did not utilise the option of scoring N4, instead scoring all potential N3 or N4 epochs as N3. Finally, panel e. shows the brain in REM sleep, the deepest sleep stage, yet showing paradoxically similar EEG traces to the wake brain, with predominant theta and slow alpha rhythms. Characteristic are the name-giving rapid eye-movements (10) and the very low muscle tone (11), only indicating heart beats here.

Sleep scoring procedure

Within the FASST toolbox, different channels were selected to represent eye movements (vertical EOGs), muscle tone (EMG right), and cortical activity from frontal (F2), central (C2), and posterior electrodes (P2, PO4), *Figure 11*. EOGs are important for detecting blinks and saccades (label 3), as indicators of an awake participant, rolling eye movements (label 6), a marker of falling asleep, and rapid eye movements (label 10), the name-giver of REM sleep stage. EMGs are similarly important for scoring periods of wakefulness and REM, with pronounced high (label 4) and low muscle tone (label 11), respectively. Posterior electrodes are helpful in identifying alpha rhythms (label 2), marking quiet alertness and thus a period of wakefulness. Sleep markers of N2, namely k complexes (label 7) and sleep spindles (label 8), originate predominantly in the central vertex (picked up by central electrodes), whilst frontal electrodes often show the highest amplitude of slow oscillations or delta waves (label 9).

Sleep scoring was initially performed independently by two scorers as follows: The data were segmented into 20-second intervals, which allowed scorers to view segments of sleep EEG signal, yet also see the interval in the context of the whole Nap phase. Looking at the entire sleep recording is an important first step, as sleep markers differ across participants, as does the dominant individual alpha frequency (e.g., Klimesch et al, 1999). Even when scoring individual segments of sleep data, it is important to consider the preceding and following epochs. The EEG signal and deduced behaviour before and after the to-be-scored time window inform and often determine the sleep stage the participant was in during the time window in question. Each segment was assigned a sleep stage ranging from awake (“0”), N1 (“1”), N2 (“2”), N3 (“3”) to

REM ("5"), creating a hypnogram of the Nap phase. Major muscle motions and arousal were marked.

The resulting hypnograms from the first and second sleep scorer were combined by a third sleep scorer. Segments of disagreement were extracted and, within the context of the entire sleep recording, re-scored by the third scorer. The final hypnogram for each participant was determined by the agreed scored segments of the first two independent scorers (agreement of 85.08% on average) and the resolved scoring by the arbitrating third scorer (remaining 14.92%).

Results

Behavioural results

Learning

Throughout the Learning task, participants detected an increasing number of targets within less search time (mean ± SEM over all conditions; Accuracy (in % of targets found) - Block 1: 95.14 ± 1.06, Block 3: 98.01 ± .65, Block 5: 99.06 ± .42; Reaction Time (search time in s) - Block 1: 7.4 ± 0.50, Block 3: 2.6 ± 0.14; Block 5: 1.6 ± 0.11). A Friedman test showed that accuracy in target detection improved significantly over blocks of trials ($X^2 = 56.19$, $p < .001$), *Figure 12c*. Moreover, a repeated-measures ANOVA on RT data revealed a significant main effect (F (1.96, 43.21) = 213.08, $p < .001$, $\eta_p^2 = .91$) and a linear contrast of block (F (1, 22) = 339.84, $p < .001$, $\eta_p^2 = .94$): Participants got progressively faster at identifying the target over blocks of learning, *Figure 12c*.

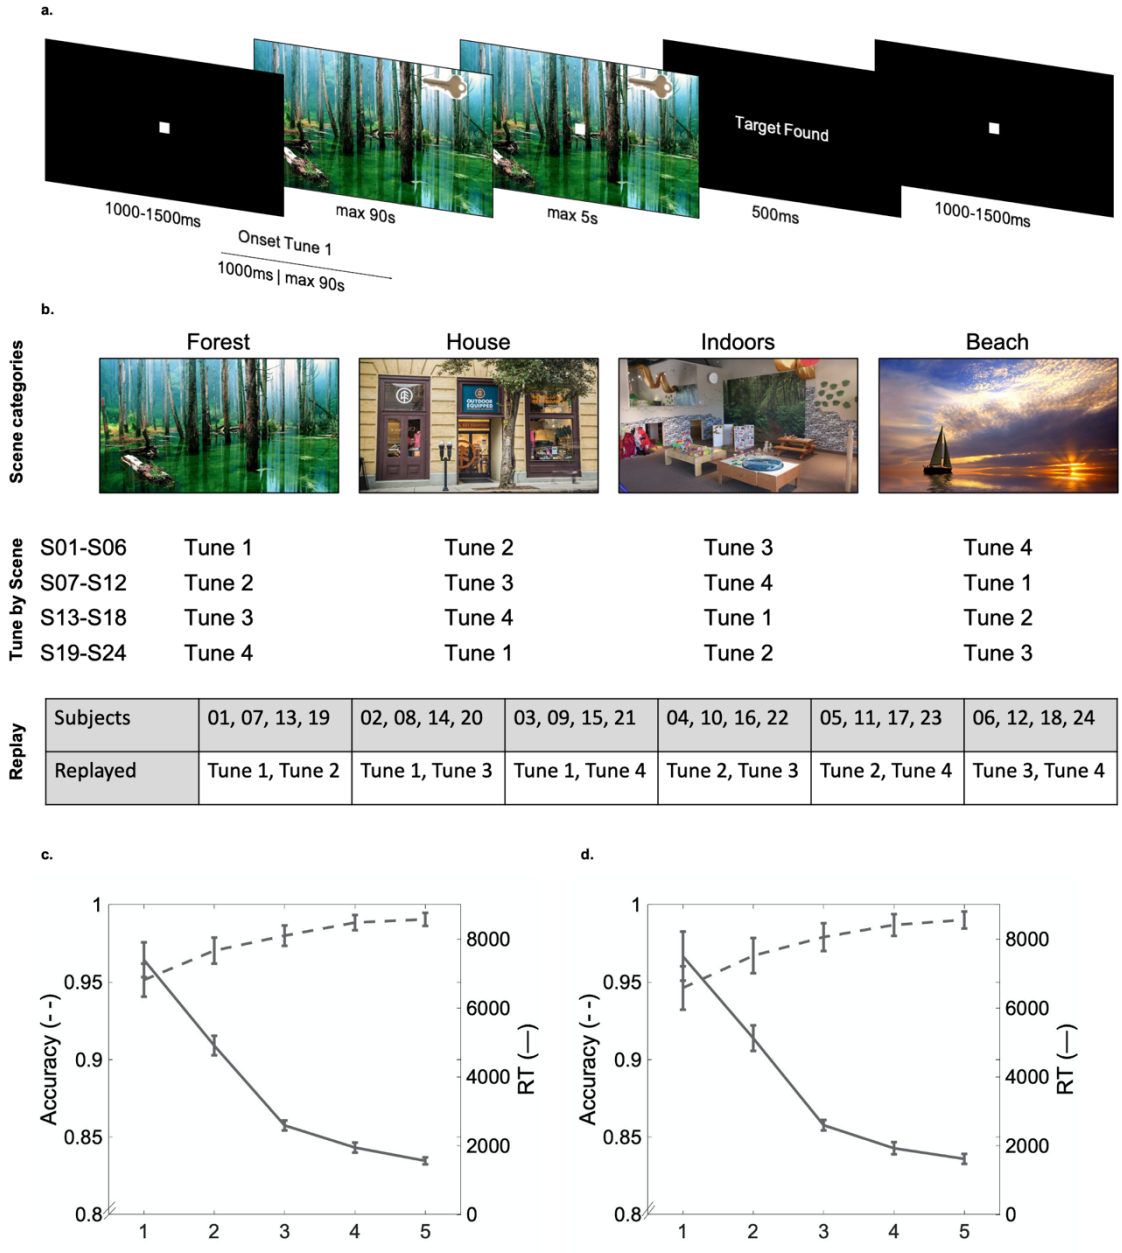


Figure 12. Panel a. provides a schematic of the Learning task: Participants were asked to search for the key in a naturalistic scene. The complex scenes were drawn from four categories, each associated with a different tune for the duration of the experiment. Stimuli size are not to scale. Panel b. illustrates the scene, tune, and replay conditions: Four scene categories and four tune categories were associated in learning. Subjects identified on the left, associated the named tune with the scene category above, balancing that each tune was equally often paired with each scene category. Of these four tunes, two were replayed during sleep, indicated by the "replay" section. This was balanced so that each tune was replayed equally often and in each combination. This meant every participant had a unique combination: For example, participant S07 had the tune-scene pairings Forest-2, House-3, Indoors-4, and Beach-1. Of these tunes, 1 and 2 were replayed during sleep. Panel c. and d. depict accuracy and RT data in response to target identification during the Learning task in all (c.) and subset of fifteen (d.) participants. Accuracy is shown to increase and RTs to decrease across blocks.

The requirement for N3-based stimulation led to the exclusion of further participants for sleep-dependent analyses. To ensure that effects were not driven by abnormal learning, I analysed the learning behaviour for the subsets of participants with sufficient sleep stimulation. In a subset of fifteen participants, with a minimum of one round of stimulation during N3, a Friedman test showed that accuracy increased significantly over blocks of trials ($X^2 = 35.03$, $p < .001$). A repeated-measures ANOVA, also confirmed a significant effect of block (F (1.85, 25.85) = 125.31, $p < .001$, $\eta_p^2 = .90$), with a linear effect showing that reaction times decreased continuously over blocks of trials (F (1, 14) = 188.05, $p < .001$, $\eta_p^2 = .93$), *Figure 12d*.

Memory

The explicit spatial recall in the pre-nap and post-nap Memory tasks confirmed that participants retained an accurate representation for target locations within at least some of the learnt scenes. Participants showed a mean spatial error of 215.01 pixels (*SE* = 20.99) before their Nap phase, and a mean spatial error of 228.86 pixels (*SE* = 21.78) after the Nap phase. Using a permutation analysis, chance performance was estimated to be 746.26 pixels (SE = 2.83). Participants were thus significantly above chance for both pre- and post-nap Memory tasks, t (22) = -12.53, $p < .001$, $d = 2.61$ and t (22) = -12.56, $p < .001$, $d = 2.62$, respectively. At a "hit criterion" of 300 pixels, these measures are equivalent to a hit rate of 79.20% (SE = 2.52) pre and of 77.99% (SE = 2.54) post nap.

Performance differed for RT but not for spatial error across Memory session, assessed on non-replayed trials to prevent confound by varying amounts of replay. Although there was a numerical increase in spatial error for post-nap (M = 235.72 pixel, *SE* = 26.61) in comparison to pre-nap memory performance (M = 225.24 pixel, *SE* =

25.49), the difference was not statistically significant (t (23) = -1.90, p = .071, d = .40), *Figure 13b.* However, there was a significant decrease in RT (t (23) = 4.98, p < .001, d = 1.04) from pre (M = 1.9s, SE = .11) to post-nap memory performance (M = 1.7s, SE = .12), *Figure 13d*.

Stimulation during sleep did not categorically modulate performance, calculated on a subset of participants with at least one round of stimulation (n=15). Two repeated-measures ANOVAs were used to calculate the effect of session as well as of sleep-based replay on spatial error and RT. The subset of participants with at least one round of stimulation showed no significant effects of session (F (1, 14) = 3.14, p = .098, η_p^2 = .18) or replay (F (1, 14) = .36, p = .557, η_p^2 = .03), nor an interaction (F (1, 14) = .34, p = .567, η_p^2 = .02) on spatial error, *Figure 13c*. With regard to reaction time, there was a main effect of session (F (1, 14) = 19.83, p = .001, η_p^2 = .59), with faster reaction times post-nap (M = 1.7s, SE = .13) in comparison to pre-nap (M = 2.0s, SE = .14), *Figure 13e*. There was no main effect of replay (F (1, 14) = .01, p = .914, η_p^2 < .01), and no interaction (F (1, 14) = 1.95, p = .185, η_p^2 = .12). Duration of sleep did not significantly covary with any of the measures reported above.

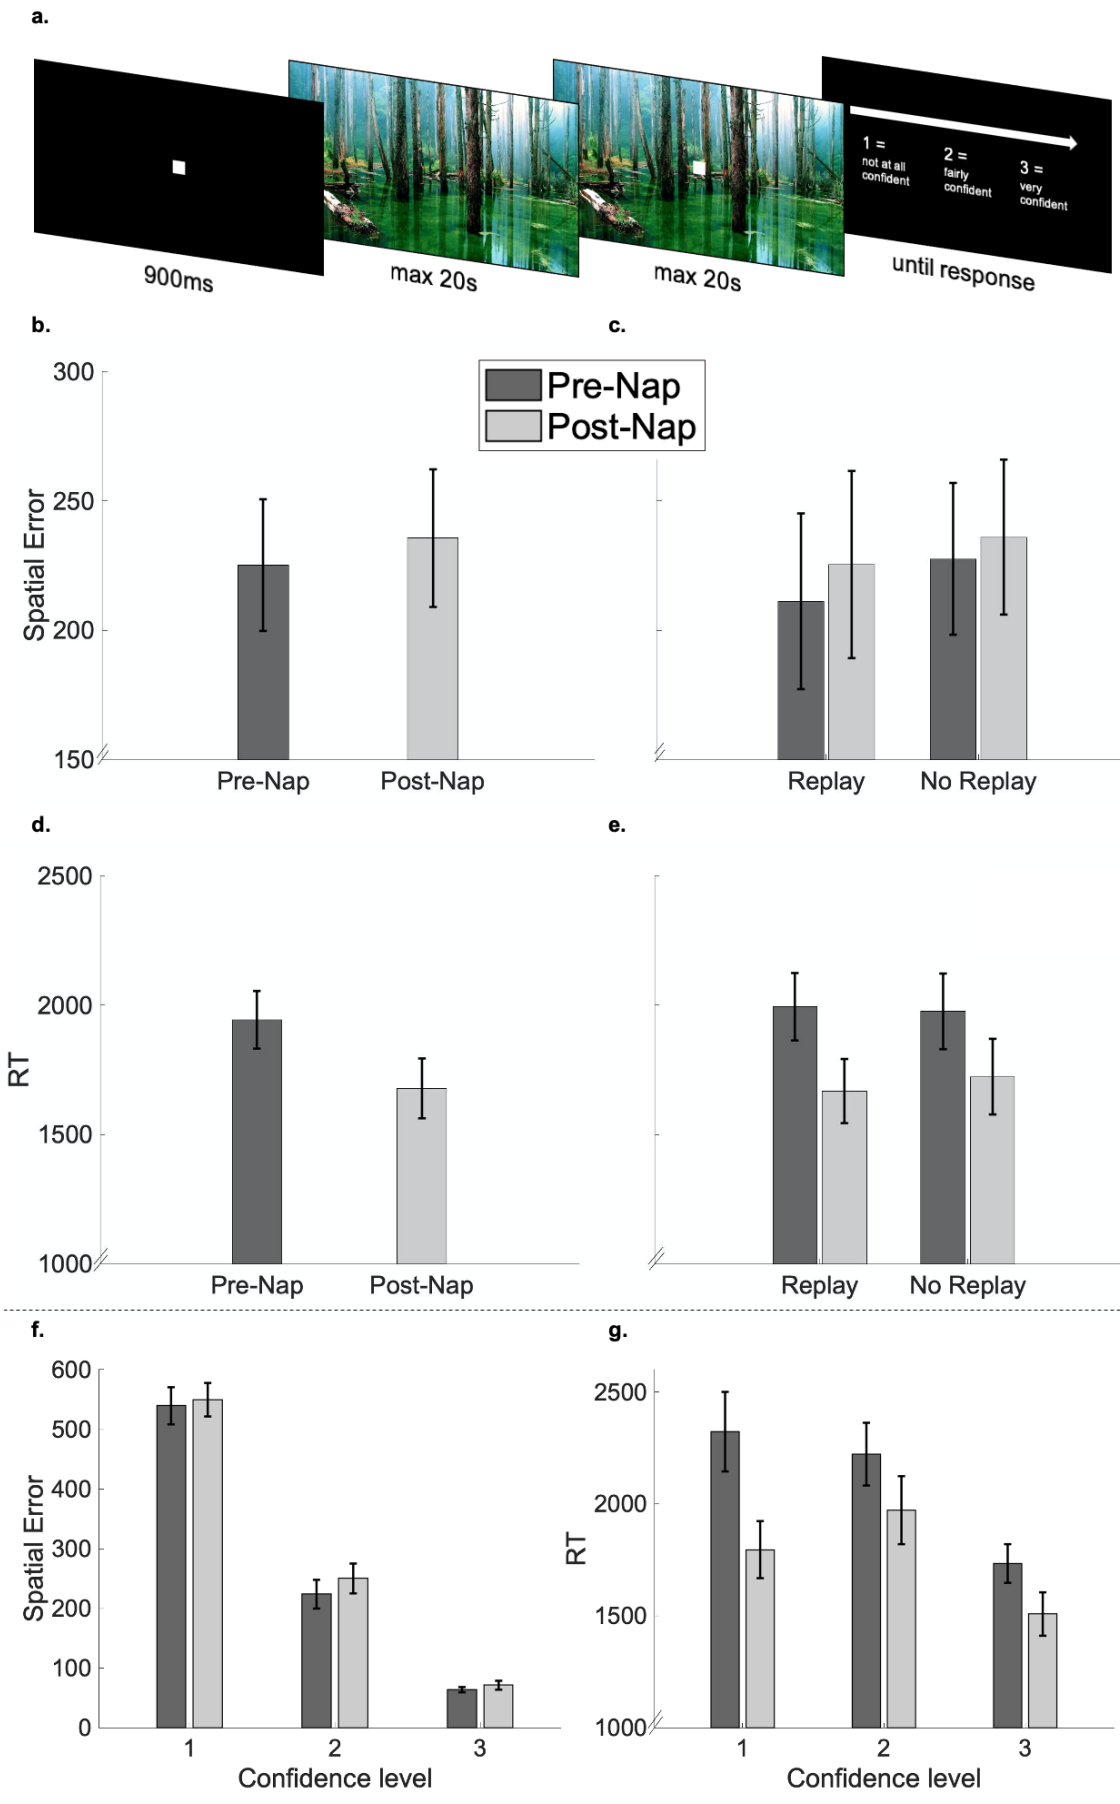

a.
900ms
max 20s
max 20s
1 = not at all confident
2 = fairly confident
3 = very confident
until response
b.
c.
Pre-Nap
Post-Nap
Spatial Error
300
250
200
150
Pre-Nap
Post-Nap
Replay
No Replay
d.
e.
RT
2500
2000
1500
1000
Pre-Nap
Post-Nap
Replay
No Replay
f.
g.
Spatial Error
600
500
400
300
200
100
0
1
2
3
Confidence level
RT
2500
2000
1500
1000
1
2
3
Confidence level

Figure 13. Panel a. is a schematic overview of the Memory task: Participants viewed the learnt scenes freely and provided information about the remembered target location, and their confidence in this memory. Stimuli sizes are not to scale. Panels b. and c. illustrate spatial error data for target placement, divided by session and replay context, for all (b.) and a subset of fifteen (c.) participants. Spatial errors did not differ significantly between pre- and post-nap, neither for all participants, nor in the subset. Panels d. and e. demonstrate reductions of RT in the post-nap Memory task, for both all (d.) and a subset of fifteen (e.) participants. No effects of replay were observed in any contrasts. The final panels show spatial error (f.) and RT data (g.) by session and confidence, demonstrating decreases in spatial error and RT with increasing confidence. Due to varying confidence ratings, the trial numbers differ for the different levels of confidence and memory session (Not, Fairly, Very (*M*, *SE*); Pre-nap: 30.96, 3.12; 42.04, 2.91; 70.96, 4.29; Post-nap: 32.74, 3.48; 37.00, 3.22; 74.26, 4.76).

As participants also reported their confidence (3-point scale) during each Memory test, I was able to assess the effects of session on spatial error and reaction time as a function of confidence. A 2 (pre or post Memory task) x 3 (confidence level) repeated-measures ANOVA showed a significant main effect of Memory session (F (1, 22) = 7.70, p = .011, η_p^2 = .26), with higher spatial errors post-nap.[3] There was also a main effect of confidence (F (2, 44) = 548.93, p < .001, η_p^2 = .96): Participants made progressively fewer spatial errors the more confident they were (linear contrast: F (1, 22) = 1059.18, p < .001, η_p^2 = .98; SE in pixels: *M*, *SE*; Confidence level 1: 544.46, 27.93; Confidence level 2: 236.92, 23.36; Confidence level 3: 67.44, 5.78), *Figure 13f*. There was no interaction between session and confidence (F (2, 44) = .93, p = .402, η_p^2 = .04), suggesting that participants had comparably accurate meta-awareness of the quality of their memory representation across sessions. The repeated-measures ANOVA on RT data resulted in significant main effects for both session (F (1, 22) = 31.56, p < .001, η_p^2 = .59) and confidence level (F (1.56, 34.23) = 24.94, p < .001, η_p^2 = .53). As discussed above, participants were faster after the nap in comparison with first Memory session. The main effect of confidence level appeared to be driven by faster responses for scenes with high confidence (RT in s: *M*, *SE*; Confidence level 1:

[3] The difference to the non-significant t-test on spatial error above was likely due to the different grouping and resulting change in weight of individual spatial error measurements.

2.06, .14; Confidence level 2: 2.10, .13; Confidence level 3: 1.62, .09). The significant interaction between session and confidence (F (2, 44) = 3.25, p = .048, η_p^2 = .13) appears to be driven by the discrepancy in RTs for low confidence scenes: Prior to the nap period, participants were slow to respond to these scenes (M = 2.32 s, SE = .18), whilst after the nap the decision time was very short (M = 1.79 s, SE = .13), especially in contrast with other blocks, *Figure 13g*. This discrepancy can also be seen in tests of within-subject contrasts, which show more pronounced effects in a linear contrast for pre-nap performance (F (1, 22) = 22.56, p < .001, η_p^2 = .51) and in a quadratic contrast for post-nap performance (F (1, 22) = 33.79, p < .001, η_p^2 = .61), respectively. However, pairwise comparisons show that RTs were significantly shorter post nap for each of the three confidence levels (1: t (22) = 5.69, p < .001, d = 1.19; 2: t (22) = 2.62, p = .016, d = .55; 3: t (22) = 5.99, p < .001, d = 1.25).

Memory-guided Attention

The Memory-guided Attention task showed that the content of LTM significantly influenced participants' detection speed for targets appearing after a scene cue. As anticipated, identification accuracy was at ceiling over all conditions: Participants responded to 96.83% (SE = 0.91) of targets and, on average, only misidentified a target in 7.88% (SE = 1.32) of catch trials. Divided by conditions, there was no significant difference in accuracy following a valid (M = 97.10%, SE = 0.83) and invalid (M = 96.56%, SE = 1.14) cue according to a Wilcoxon Signed-Rank test (V = 62.00, p = .932), with both conditions showing a median of 97.92%. The following analyses on RT data only used accurate trials (number of trials used: M = 92.96, SE = .87). A paired samples t-test showed a significant effect of cue condition on RT data (t (22) = -6.90, p < .001, d = 1.44), showing faster reaction times after valid (M = 368.06 ms, SE = 11.51)

than after invalid cue (M = 408.02 ms, SE = 16.11), *Figure 14b*. The preceding analyses were using only non-replayed trials, unconfounded by potential effects of stimulation. The following analyses assessed the effects of replay and were thus run on a subset of participants who received at least one round of stimulation during sleep.

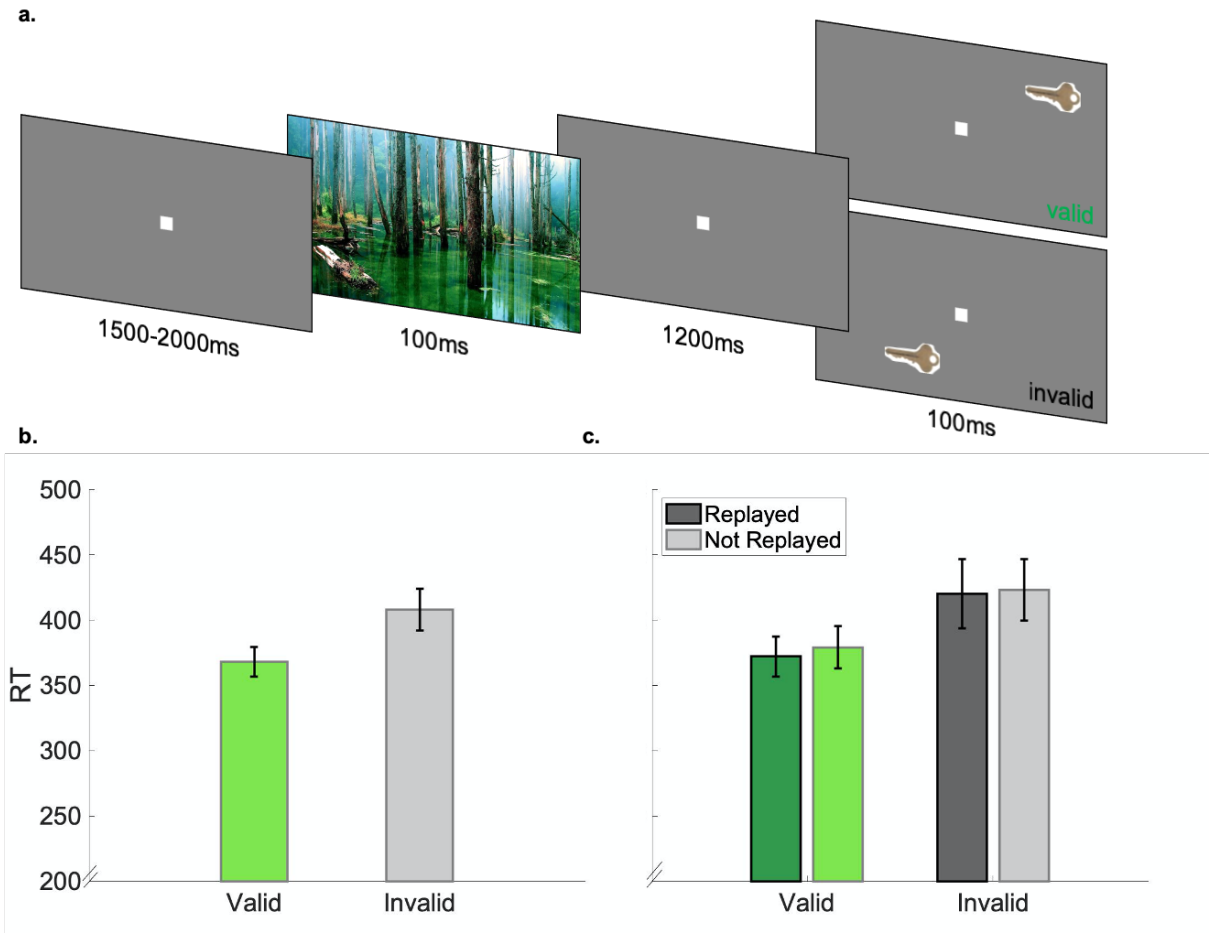


Figure 14. Panel a. provides a schematic overview of the Memory-guided Attention task: Learnt scenes without key represented memory cues to orient attention towards the remembered key location. Half of the keys flashed up in this "valid" location, half were presented at other "invalid" locations. Half of each valid and invalid targets were association with sleep-based replay. Stimuli sizes are not to scale. Panel b. illustrates RT data in all participants by cue condition (valid or invalid) for trials without replay (lighter colours, grey border): Valid trials were associated with faster reaction times. Panel c. demonstrates that RTs in a subset of participants with at least one round of stimulation during sleep were significantly impacted by cue condition: Valid trials were associated with significantly decreased response times. Replayed trials (darker colours, dark border) did not differ significantly from non-replayed trials (lighter colours, grey border).

Wilcoxon Signed-Rank tests on accuracy data in the subset with the minimum of one round of stimulation (n=15) showed no effects of validity for replayed (V = 11.00,

p = .418) or non-replayed trials (V = 7.00, p = .263). There was also no effect for replay within the cue condition, showing no significant effect for scenes after valid cues (V = 15.50, p = .863) or invalid cues (V = 5.50, p = .168). A repeated-measures ANOVA on the RT data emphasised the influence of cue condition, showing a significant main effect of cue (F (1, 14) = 23.89, p < .001, η_p^2 = .63): Participants were faster in trials with valid cues (M = 375.72 ms, SE = 15.57) than with invalid cues (M = 421.75 ms, SE = 24.74), *Figure 14c*. However, there was no main effect of replay (F (1, 14) = 1.56, p = .232, η_p^2 = .10), with participants showing similar reaction times to replayed trials (M = 396.35 ms, SE = 20.28) and not replayed ones (M = 401.12 ms, SE = 19.59). There was no interaction between validity and replay (F (1, 14) = .09, p = .763, η_p^2 = .01).

Further, I assessed whether the cueing effect was associated with memory performance. I correlated spatial error in the pre- and post-nap Memory task using non-replayed trials with a normalised validity score. The correlations across tasks were not significant for pre nap (r_s = -.26, p = .221), *Figure* 15a, or post-nap (r_s = -.33, p = .125), *Figure 15b*.

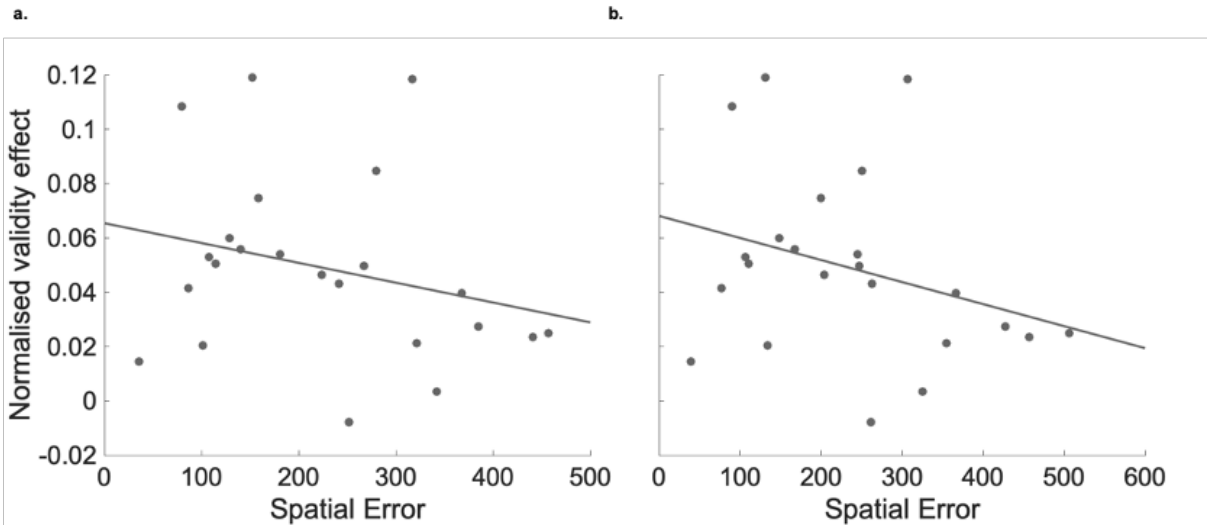


Figure 15. The correlation between Memory and Memory-guided Attention task performance was not significant for pre-nap (panel a.) or post-nap performance (panel b.). The direction would have indicated that lower spatial errors were associated with larger RT benefit in response to valid as opposed to invalid memory cues.

Morningness-Eveningness Questionnaire results

The results of the Morningness-Eveningness Questionnaire indicated that there was no effect of individual types on sleep and behavioural measures. Fourteen participants scored as intermediate type, with only five participants scoring each as evening types (four moderate, one definite) and as morning types (four moderate, one definite). Participants' score on the questionnaire did not correlate significantly with measures of sleep (sleep duration: r_s = .23, p = .282; N3 duration: r_s = -.07, p = .760) or relevant behavioural measures (post-nap memory performance: r_s = .25, p = .254; normalised forgetting: r_s = -.14, p = .516; normalised validity: r_s = .15, p = .508; differential cueing effect: r_s = -.15, p = .483).

Sleep-dependent results

Sleep physiological measures across participants are summarised in Table 1. There was a high level of agreement across the two individual scorers (M = 85.08%, SE = 1.32) and with the resulting composite score, M = 88.10% (SE = 1.04) and M = 96.51% (SE = .58), respectively. The agreement of N3 scoring, the most relevant sleep stage for this study, was very high: The first and second scorer showed an overlap with the final score of M = 94.19% (SE = 2.17) and M = 99.61% (SE = 0.35) for N3.

All participants were able to sleep within the Nap phase, as evidenced by the high percentage of the session spent in sleep (M = 77.91%, SE = 2.75). The quality of the sleep may have been a bit interrupted, or at least less stable than sleep period during an overnight sleep, as shown by the high number of state transitions (M = 27.04, SE = 2.45) during the relatively short window of potential sleep (M = 83.78 minutes, SE = 1.18). The majority of the sleep (47.69%) was spent in N2. The second most common sleep stage was N3 (27.47%), considered to be essential for memory consolidation and

target for sleep-based stimulation. The transitional stage N1 accounted for around a fifth of the sleeping time (20.51%), and there was very little time in REM sleep (4.34%). It is also notable that only seven of the twenty-three participants entered into REM sleep, whilst nineteen participants experienced at least some N3 sleep. Stimulation was mainly intended to occur in N3, and sleep physiology confirmed that this was successful, with 71.39% of stimulation occurring during N3. However, importantly, only fifteen participants received at least one round of stimulation during N3, constituting a subset for analysis, and only twelve participants received at least five rounds.

Table 1. Measures are shown over all subjects, and over a subset of 15 participants that received a minimum of one round of stimulation during SWS.

Sleep Statistics	**Mean +/- SEM (n = 23)**	**Mean +/- SEM (n = 15)**
Sleep efficiency (% of session spent asleep)	77.91% (2.75)	79.50% (2.72)
Number of arousals and awakenings	2.78 (0.48)	2.73 (0.55)
Number of state transitions	27.04 (2.45)	22.00 (2.05)
Sleep onset latency (in min)	9.55 (1.09)	10.02 (1.39)
Duration of sleep session (in min)	83.78 (1.18)	82.76 (1.56)
N1 (in min)	13.35 (1.32)	11.18 (1.59)
N2 (in min)	31.04 (3.05)	26.97 (3.73)
N3 (in min)	17.88 (2.75)	24.49 (2.80)
REM (in min)	2.83 (0.96)	3.24 (1.30)
Rounds of stimulation in SWS	10.96 (2.20)	16.80 (2.19)
Rounds of stimulation in sleep	15.35 (2.33)	21.67 (2.16)

The subset of fifteen participants showed similar overall sleep physiology through the length of the Nap phase (M = 82.76 minutes, SE = 1.56), spending the majority of the duration in some stage of sleep (M = 79.50%, SE = 2.72), *Table 1*. They showed less, albeit still high number of state transitions (M = 22.00, SE = 2.05). Whilst N2 still accounted for the majority of sleep (40.94%), this participant subgroup showed a much higher percentage of sleep in N3 (37.17%). This was most likely a by-product of the selection criteria for this subgroup, as a minimum in stimulation during N3 leads to an emphasis on N3 duration. The remainder of sleep was scored as N1 (16.97%) or REM sleep (4.92%). Again, only a fraction of participants, namely five, entered into REM sleep during the Nap phase. On average, participants received 16.80 (SE = 2.19) rounds of stimulation during N3, equivalent to 77.54% of all received stimulation.

Effects of sleep on Memory

A Spearman's rank correlation test confirmed a significant positive correlation of overall sleep duration with spatial error in the post-nap Memory task: Longer sleep duration was associated with a higher spatial error (r_s = .51, p = .013) across all participants, *Figure 16a*. Sleep duration was also positively correlated with spatial error in the pre-nap Memory session (r_s = .53, p = .010), as one might expect given the high correlation of pre- and post-nap Memory performance (r_s = .98, p < .001). There were no significant correlations of the post-nap spatial error in the Memory task with the duration of sleep stages (N1: r_s = .01, p = .971; N2: r_s = .24, p = .397; N3: r_s = .10, p = .650), *Figure 16b*. There was no correlation of normalised forgetting, the change in memory from pre- to post-nap session, with overall sleep duration (r_s < .01, p = .984). Further, no correlations of normalised forgetting with individual sleep stages were observed (N1: r_s = -.38, p = .076; N2: r_s = .14, p = .221; N3: r_s = .04, p = .847).

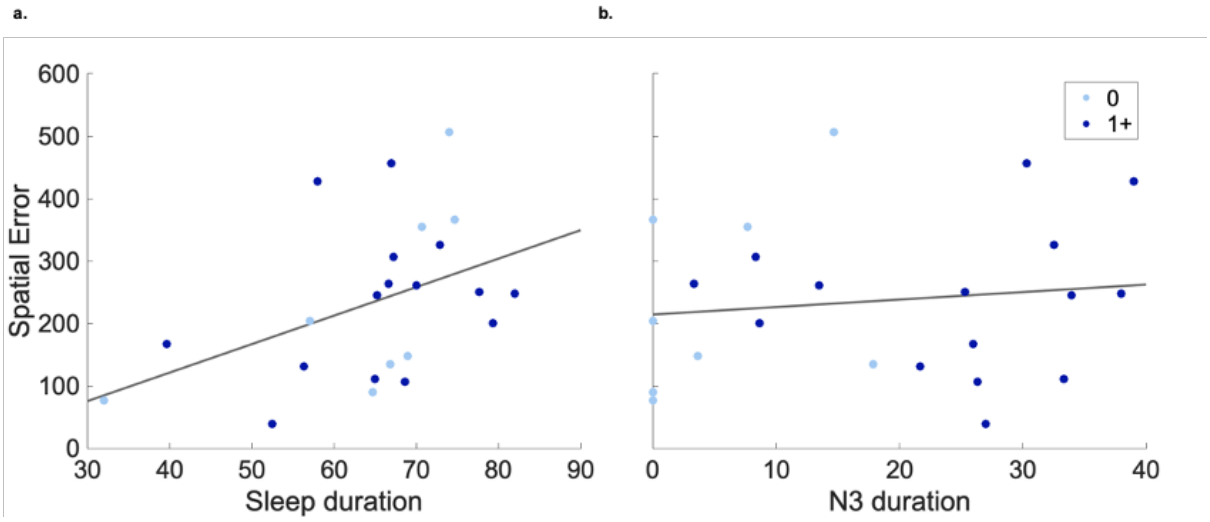


Figure 16. Correlations are shown between spatial error performance in the post-nap Memory task in scenes not associated with replay and sleep durations. Panel a. shows a significant positive relationship with overall sleep duration. Panels b. illustrates the absence of significant correlations with duration in the sleep stage of most interest, N3. Colour of the dots refers to the rounds of stimulation the respective participant received.

Effect of sleep on Memory-guided Attention

The relationship of sleep physiology and performance in the Memory-guided Attention task was estimated by correlating sleep duration with a normalised validity effect. A Spearman's rank correlation of normalised validity with overall sleep duration was not significant (r_s = .01, p = .960), *Figure 17a*. Further, the normalised validity effect was not significantly correlated with the duration of any of the NREM sleep stages, namely N1 (r_s = -.20, p = .371), N2 (r_s = .24, p = .288), or N3 (r_s = -.18, p = .433), *Figure 17b*.

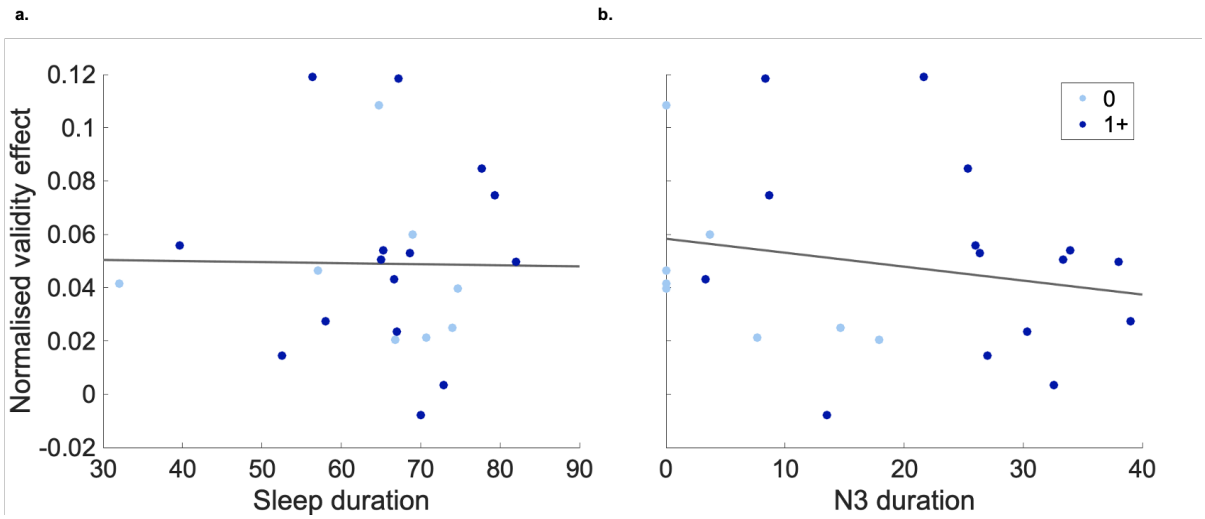


Figure 17. Correlations shown are shown between a normalised validity effect, calculated based on non-replayed scenes, and durations of overall sleep (a.) and the most relevant sleep stage N3 (b.). No significant correlations were observed. Colour of the dots refers to the rounds of stimulation the respective participant received.

Stimulation-dependent results

The behavioural results of the Memory task and the Memory-guided Attention task were also analysed with respect to stimulation during sleep. Considering the requirement for stimulation for these analyses, the data for the following sections were drawn from trials associated with replay and from subjects who had received a minimum of one round of stimulation in slow-wave sleep (n=15). Exploratory analyses were performed to gage the effect of increased stimulation on a subset of participants with at least five rounds of stimulation in slow-wave sleep (n=12).

Effects of stimulation on Memory

I did not observe performance modulation in the post-nap Memory task based on at least a single round of stimulation during slow-wave sleep. There was no significant correlation of stimulation duration with spatial error in scenes associated with retrieval, (n=15: r_s = .16, p = .567), *Figure 18a*, nor with the normalised forgetting score (n=15: r_s = -.30, p = .280), *Figure 18b*.

It was noteworthy that in the further subset with higher number of stimulations, stronger relationships emerged: The correlation between stimulation duration and spatial error was not significant, despite the numerical increase of the correlation coefficient (n=12: r_s = .50, p = .097), *Figure 18a*. In contrast, the correlation with the normalised forgetting score was significant (n=12: r_s = -.65, p = .023), *Figure 18b*. There was no effect of differential forgetting in either group (n=15: r_s = .05, p = .850; n=12: r_s = -.30, p = .336).

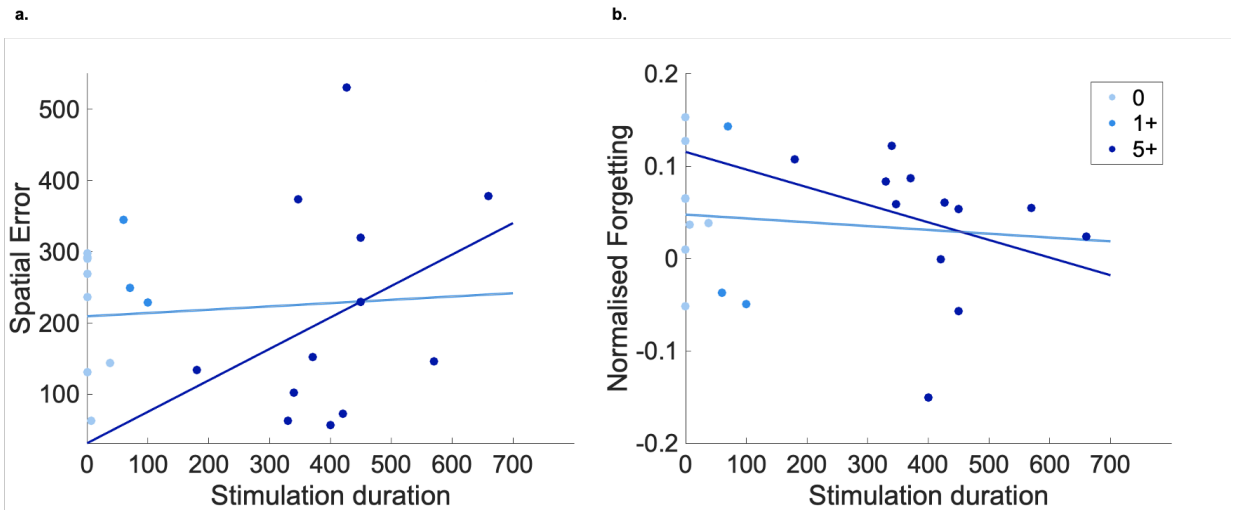


Figure 18. Panel a. shows the relationship of memory performance as spatial error in Euclidean distance with the stimulation duration in seconds. Colour of the dots refers to the rounds of stimulation the respective participant received. Darker colours were associated with at least one round of stimulation, with the dark blue colour indicating at least five rounds of stimulation. This differentiation resulted in the colour separation along the x axis. Memory performance in replay-associated trials showed no significant correlation with stimulation duration for at least one (medium blue medium and dark blue, n=15) or five rounds of stimulation (dark blue, n=12). Though not significant, the direction of the regression line for the five-rounds subset would indicate an association of larger spatial error with longer stimulation. Panel b. illustrates the association of stimulation with normalised forgetting from pre-to post-nap Memory session, with larger scores indicating increased forgetting. There was no significant correlation over participants with at least one round of stimulation (medium blue medium and dark blue, n=15), yet a significant negative relationship emerged for participants with five rounds of stimulation (dark blue, n=12): This would point towards less forgetting, less decline in spatial precision, across Memory sessions with increased stimulation.

Effects of stimulation on Memory-guided Attention

I assessed the relationship of Memory-guided Attention task performance with the sleep-based stimulation by correlating it with the normalised validity effect in replay scenes. Surprisingly, I found a negative correlation (r_s = -.59, p = .028), suggesting that more stimulation was associated with less benefit from a valid memory cue in comparison with an invalid one, *Figure 19a*. This relationship was preserved in the exploratory analysis with participants who received at least five rounds of stimulation (r_s = -.64, p = .033): Again, the correlation showed a negative relationship, indicating a detrimental effect of stimulation duration on the normalised validity effect.

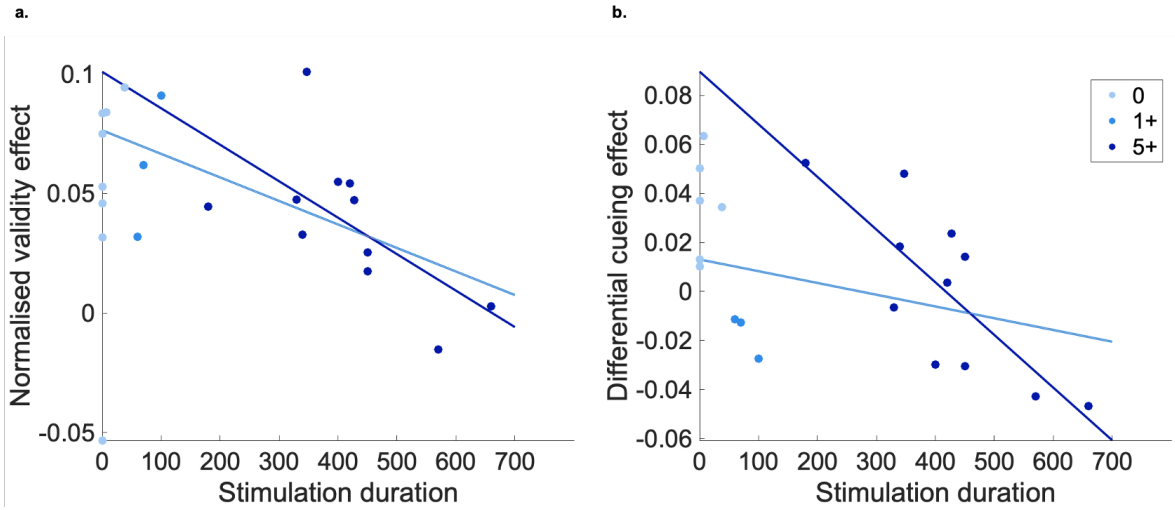


Figure 19. Panel a. The stimulation duration during sleep per participant is plotted against the normalised validity effect, calculated from replayed scenes. Colour of the dots refers to the rounds of stimulation the respective participant received. Both regression lines, based on a minimum of one (medium and dark blue, n=15) and five (dark blue only, n=12) rounds of stimulation, show a significant decrease in validity effect with increasing stimulation. Stimulation was associated with diminished utility of the replayed scenes as memory cues. Panel b. captures the differential effect of memory cues associated with and without replay, by the amount of stimulation received. There was no significant correlation in participants with a minimum of one round of stimulation (medium blue medium and dark blue, n=15). However, the strongly negative regression line (dark blue, n=12) illustrates the negative correlation of stimulation duration with differential cueing effect for participants with at least five rounds of stimulation: The more stimulation participants received, the more the cueing benefit was pronounced in non-replayed scenes as opposed to replayed scenes.

Exploring this effect further, I evaluated the relationship of stimulation duration on the differential cueing effect. The differential cueing effect was calculated by subtracting the normalised validity effect for non-replayed trials from the normalised validity effect for replayed trials, conveying an insight in the differential effect of replay on the use of memory to guide attention. There was no significant correlation for the subsection of participants with a minimum of one stimulation round (r_s = -.34, p = .230). However, there was a strong negative correlation in participants with more stimulation (r_s = -.69, p = .018), *Figure 19b*.

Discussion

The present set of results indicated that the learning task was successful at building robust long-term memories for the location of targets within scenes and that such memories effectively drove spatial attention in a Memory-guided Attention task. Yet, I was unable to detect benefits of sleep-dependent consolidation for the quality of memories nor their utility in driving attention. Whilst there may have been weak indications of a protective role of sleep-based stimulation from forgetting, I could not find broad behavioural benefits in memory retention. Surprisingly, I found detrimental effects of stimulation during sleep on the utility of memories as attentional guides.

This study replicated and extended main behavioural findings from similar memory-guided attention paradigms. Participants were able to form location and scene associations progressively over the learning blocks, showing speeded target detection with increasing accuracy. The Learning task was thus successful irrespective of the differentiation in scene categories and the associated auditory stimuli. Learning performance was stable across participants, showing an equivalent pattern of results for

a subset of participants, who received a minimum amount of stimulation during slow-wave sleep.

Participants successfully formed robust memories, which persisted over multiple delays. Notably, one of these delays included a Nap phase. Overall, participants showed no measurable forgetting across this second delay with regards to spatial precision of the location memory. However, participants were faster to respond in the post-nap Memory session. Whilst this could be interpreted in the light of a speed-accuracy trade-off, it definitely indicates that participants did not show sleep inertia, which would have prevented them from performing the cognitive tasks post-nap to the same level. However, any other circadian effects might be difficult to take into account, as all participants by design completed tasks at the same time of the day.

In general, participants exhibited good meta-awareness of their own performance: with increasing confidence, participants also displayed higher spatial precision at a greater response speed in the memory. This meta-awareness did not change across Memory session with respect to spatial precision. However, an interesting interaction emerged with regard to the speed with which participants responded. Within the first Memory session, the relationship between response time and confidence was straightforward and linear, showing decreased response times with increasing confidence. Post-nap, participants showed shorter reaction times for both extremes, where they were sure of either remembering or not remembering the scene and associated location. On the contrary, scenes with intermediary confidence appeared to warrant longer consideration.

Memories were successful in guiding attention, irrespective of replay and associated auditory context. In agreement with results from previous simple-detection versions of memory-guided attention tasks (e.g., Summerfield et al., 2006), participants

performed this task at a very high level of accuracy across conditions, as can be expected with the comparably low task difficulty. The reaction-time measure is more sensitive, and hence appropriate, in differentiating the effect of the memory cue on performance. Indeed, participants were faster in response to an appearing target, if the target appeared in the location predicted by the memory cue. The memory cueing benefit was stable across participants. It was observed across all participants as well as in the subset with a minimum amount of sleep-based stimulation. Within this study, the benefit conferred by valid memory cues grew numerically as spatial errors in memory diminished, but the correlations did not reach significance for errors in either Memory session. The lack of correlation differentiation might be explained by the low number of participants as well as by the observation in the literature that memories that drive the effects of memory-guided attention can differ from the memories that drive explicit memory performance (Salvato et al, 2016a, 2016b).

The novelty of this memory-guided attention was the sleep component. The Nap phase was successful: the majority of participants not only fell asleep but also entered deep sleep stages such as N2 and N3 for extended periods of time. This was confirmed by three independent scorers, showing a high degree of agreement across sleep scores. The setup of the Nap phase was thus successful, allowing sleep-based auditory stimulation unbeknownst to the sleeping participants. The timing of the Nap phase, aiming for the natural circadian dip in the early afternoon, appeared successful. Despite being drawn from a spread of sleep types (morning or evening types) with varying sleeping habits participants fell asleep to equivalent extents.

A stated aim of this study was to address whether memories used to guide attention are subject to sleep-dependent consolidation. At least within the constraints of this study, I did not find evidence to support my hypothesis: Neither the rate of

forgetting across Memory sessions, nor strength of the cueing benefit was influenced by the time participants spent in sleep or in specific sleep stages such as slow-wave sleep.

However, and interestingly, the duration spent in sleep was associated with the spatial error in the post-nap Memory session. Contrasting to literature findings, longer sleep was associated with higher average memory error. It is worth noting that the normalised forgetting measure, showing non-significant results, is a more sensitive measure as it takes out some of the inter-person variance by comparing participants' performance across sessions and normalising the resulting score. However, it is still worth discussing why longer sleep was correlated with larger spatial memory errors. The observations of a similar correlation for pre-nap memory performance and of the high correlation of pre- and post-nap memory performance suggest that it is unlikely that the duration in sleep had a differential impact on post-nap memory error. Instead, the direction of the correlation could be an artefact of the design: It might be as simple as that participants who were more tired and thus potentially performed worse on the Memory task ended up sleeping longer. Despite the attempts to recruit participants from a broad background, the location of the study made it very likely that the majority of participants were students, amongst which irregular sleep habits and sleep deprivation might be more common than in the general public. Yet, of course, this is just speculation. In any case, any negative association of sleep duration with memory performance did not influence the use of memories in the guidance of attention.

Potentially there would have been stronger effects of sleep, especially slow-wave sleep, with a different task set-up. Within the current study, participants were not explicitly aware that they would be asked to recall the information later. However, there are findings in the literature supporting the idea that the impact of slow-wave sleep on

memory consolidation depends on future relevance: Wilhelm and colleagues (2011b) observed that recall and slow-wave sleep only correlated if the memory recall was not a surprise. Perhaps the results of the present study were influenced by the fact that participants did not necessarily expect a memory recall task after their Nap phase. Further, it has been observed previously that profit from sleep-dependent consolidation appears to be greater for information that was more weakly encoded (Drosopoulos et al, 2007). More nuanced effects of sleep on memory, and thus memory-guided attention, might be more effectively revealed in fine-tuned analyses based on trial-by-trial measures. Given the constraints of trials, and participants, I would not have the necessary power to give these analyses credibility. Future studies, for example using online learning and the potential of over-training participants on the location-scene-tune associations, would allow using a larger number of trials and a bigger sample size.

Auditory stimulation during slow-wave sleep did not influence memory retrieval and precision according to the associated scene category, nor the utility of memories in guiding attention. Within the subset of participants who received a minimum of one round of stimulation during slow-wave sleep, I did not observe a modulation of spatial error or reaction time based on whether the scene was from a category associated with a replayed tune. Consequently, my hypothesis about the efficacy of abstract, conceptually unrelated auditory cues in biasing memory consolidation was not unambiguously supported by the current set of results. There was also no categorical effect of replay on memory-guided attention: Reaction times to memory cues did not differ significantly depending on association with replay. Somewhat unsurprisingly, given the absence of a clear effect of targeted memory reactivation on memory performance, I also did not find apparent evidence supporting the idea that sleep-based stimulation affects the use of memories in guiding attention.

There are multiple possible explanations for the lack of a clear, typical targeted-memory reactivation effect in this study. Firstly, by the nature of the task design, there was a difference in the tune played during the Learning task and the nap period, for example with reference to the duration of the auditory tunes and the overlay of brown noise and the tunes. This differential pattern could have been exaggerated by the fading in of the tunes during sleep, which – whilst necessary to minimise arousal risk – would have somewhat masked the beginning of the tune. However, the beginning of the tune may have been the only part of the tune participants were reliably and repeatedly exposed to during the Learning task: Especially within the later blocks, reaction times were very fast, and thus participants only heard the first few seconds of the tune. This discrepancy of the auditory presentation between the Learning and nap component could have contributed to the lack of a categorical effect of targeted memory reactivation.

Secondly, a distinctive particularity with the auditory stimuli in this study is that they were not intrinsically linked to the associated location-scene pairings beyond their co-occurrence during encoding. Many auditory targeted memory reactivation studies have used sounds that were conceptually related to the target objects, such as a meow with a cat (e.g., Oudiette et al, 2013). This had potential implications during hippocampal reactivation, on whether the sounds triggered concepts to be reactivated first rather than the novel mnemonic association. For this study, the tunes were newly composed, so had no prior associations, and no conceptual link with the scene-location pairings beyond their co-occurrence during the learning phase of the study. Notably, the question on whether auditory tunes need to be conceptually related for effective targeted memory reactivation is still debated in the literature (Schouten et al, 2017).

Thirdly, given that auditory and visual stimuli had no intrinsic link, the association might need to be "over-trained" to cause reliable targeted reactivation during sleep. The few exposures during the Learning paradigm may not have been enough to create a sufficiently strong association. Indeed, comparable studies with no intrinsic association of sound and task included more repetitions of the associations and specific task subsets to strengthen the associations (Schechtman et al, 2020). Further, the number of exposures to stimuli might not only be a factor during the learning, but also for the sleep-based stimulation itself. The reported frequency of stimulation during sleep is much higher in studies that find effects of targeted memory reactivation (e.g., Schechtman et al, 2020, around 236 instances of cueing during sleep). An increased number of stimulations might be especially important within a study like the present one, given the use of much longer and more complex auditory stimuli. The complexity of stimuli by itself should not have prevented targeted memory reactivation, as complex auditory stimuli have been used successfully before (for example, Antony et al, 2012). Instead, the number of sleep-based repetitions might be decisive for the occurrence of consolidation.

Indeed, I did find evidence for the idea that the effect of targeted memory reactivation might be contingent on the quantity of sleep-based stimulation. Taking the amount of stimulation into account allowed more nuanced insights in potential effects of stimulation on behavioural performance in this paradigm. I did not observe a correlation between stimulation duration and memory performance over participants with a minimum of one round of stimulation. However, exploratory analyses suggest a relationship emerges with more rounds of stimulation: Longer stimulation duration was associated with less normalised forgetting, thus a smaller decline in spatial precision across Memory sessions. These results are more in line with my hypothesis and would

replicate the literature on targeted memory reactivations. Stimulation appeared to have a beneficial effect on the retention of memory, at least with regard to spatial precision. Given that these analyses were exploratory and based on a small sample size, it is important not to overinterpret these findings. As outlined above, these analyses would also benefit from a larger sample size, and potentially a larger number of trials, achieved for example by an online training regimen.

Strikingly, and contrary to my predictions, I found a strong negative correlation of stimulation duration and the validity effect in the Memory-guided Attention task: Participants gained less from a valid memory cue in comparison with invalid cues, the more stimulation they had received. Longer stimulation duration was also associated with less benefit of memory-guided attention for replayed scenes as opposed to non-replayed ones. These results were counter-intuitive: The attempt to cue memory consolidation with targeted memory reactivations actually appeared to decrease the utility of associated memories as attentional guides. This negative relationship warrants reflection.

On the one hand, it is possible that there may be an inherent dissociation in the types of memories driving attentional orienting vs. explicit recall, as mentioned above (Salvato et al, 2016a, 2016b). Memories relevant for explicit recall might benefit from targeted memory reactivation, as demonstrated by the literature discussed above (see Hu et al, 2020, for meta-analysis). By contrast, memories implicitly driving the quick guidance of spatial attention, possibly remaining below conscious awareness, might be influenced by reactivation in a way that hinders their use in attentional guidance. Further studies with larger power would be needed though to support such a claim.

On the other hand, or additionally, the detrimental effect of sleep-based stimulation on memory-guided attention might be a consequence of specific choices in

the task design. Within this study, each tune was associated with a whole category of stimuli. It is conceivable, therefore, that if targeted memory reactivation occurred, it might have resulted in reactivating multiple memory traces at once. The concurrently activated memories, placed in a fragile state, may therefore have interfered with one another and disrupted individual location associations in other scenes of that category. Furthermore, interference on such a level might be enough to affect implicit use of memories negatively, yet leave explicit recall, such as in the Memory task, unharmed.

In general, there are different aspects of the set-up of this sleep study which may have influenced the success of targeted memory reactivations and putative corresponding consolidation. Some of these effects are methodological, some based on the nature of sleep architecture and function of different sleep stages. Addressing these issues systematically will improve and advance the research methodologies and lead to our better understanding of the role of different sleep stages on memory quality and utilisation. Future studies may thus be able to optimise consolidation effects and minimise or even overcome certain confounds altogether.

Firstly, there is always a chance of arousal following auditory targeted memory reactivation. Arousal following stimulus presentation can disturb memory consolidation (Göldi and Rasch, 2019), and thus the efficacy of targeted memory reactivation. Within the present set-up, this could not be very well controlled, as noise level could not be adjusted flexibly in response to observed arousals. Tools are being developed to adjust the noise level of the auditory stimuli flexibly to stay audible yet below the personal arousal threshold, which obviously varies depending on sleep stage.

Secondly, within the targeted memory reactivation literature, there is some discussion about the effect of different oscillatory properties at the time of stimulation. The phase of the slow oscillations during stimulation has been observed to be important

(Batterink et al, 2016), usually with beneficial effects seen when stimulating on the up phase of the slow oscillation (Ong et al, 2016; Göldi et al, 2019). Different methodologies are being developed to control the timing of the stimulation relative to slow oscillation phase (for example, Santostasi et al, 2016). Other results stress the importance of avoiding the spindle refractory period, a period directly after a spindle occurred, for optimal stimulation conditions (Antony et al, 2018). Further, there is a need for more research to determine the optimal timing of stimulation for complex stimuli: Does the beginning of the stimulation have to coincide with optimal slow oscillation phase and spindle state? Or rather the most recognisable or defining aspect?

A third general consideration is the function of the sleep stages themselves and their relation to the design of the paradigm. There are numerous studies reporting the effect of targeted memory reactivations during afternoon naps, focusing solely on slow-wave sleep (for example, Antony et al, 2012; Ong et al, 2016). Others argue that REM sleep might serve an important complementary function in memory consolidation: Whilst targeted memory reactivation in slow-wave sleep might "tag" memory for further consolidation, REM sleep might be important for the actual integration (Tamminen et al, 2017). Tamminen and colleagues found that despite targeting slow-wave sleep with stimulation, the triggered integration of new memories was actually predicted by the duration of REM sleep. The question on whether slow-wave sleep is sufficient for consolidation, or whether it relies on subsequent REM sleep-based processes, is still a matter of debate in the field. Importantly with regard to the setup of a sleep study, REM sleep is not often present in afternoon naps, as mirrored in the present study. However, if REM sleep is shown to have a mediating, or optimising influence on memory consolidation triggered by targeted memory reactivation, future paradigms should reflect this. One easy preliminary way to address this line of enquiry

is by retesting participants the next day, as overnight sleep, in all probability including REM sleep, might give rise to more pronounced effects of the sleep-based stimulation the day before.

In conclusion, I have shown that participants formed and retained memories for location and scene associations within auditory contexts. They were able to use these memories to guide their spatial attention effectively. Within this study, participants' memory did not benefit overall from sleep or sleep-based stimulation, neither did their ability to use the memories as attentional guides. However, I did observe modulations in behaviour based on the quantity of stimulation received: Increased stimulation may have protected the spatial precision of the memories yet impacted the utility of the memory negatively with respect to attentional guidance. Potential reasons for these findings were discussed.

Chapter IV. Selective attention during retrieval

Introduction

The previous chapters have addressed the relationship between attention and memory at encoding (Chapter II) and during consolidation (Chapter III). In this chapter I examine the relationship between attention and memory at retrieval.

Even though we ourselves might not be fully aware of it, most acts of our memory retrieval are selective. When we attempt to recall a certain holiday experience or try to remember the name of a long-lost acquaintance, we search for this episode or even this one item among countless other episodes and items. There is simply too much information to be considered all at once; instead processing needs to be selective (Chun et al, 2011).

Some evidence for the idea of competing memories comes from neuroimaging studies. Memory retrieval has been shown to involve the reactivation of neural regions recruited during encoding (for review see Danker & Anderson, 2010). One might therefore expect that memoranda with similar patterns of activation during encoding to compete. Indeed, Kuhl and colleagues (2011, 2012) observed overlapping and competing brain activity in ventral temporal cortex during the retrieval of memories with similar content. Their findings further suggest a functional link between competition between memories and retrieval success; stronger reactivations of competing memories led to more retrieval errors. It has been proposed that the process of selecting content during memory retrieval may lead to suppression of competing content and strengthening of the neural representation of selected content over time (Wimber et al, 2015). Further evidence for the selectivity of memory retrieval comes from studies into distraction (e.g., Wais et al, 2010). Tarder-Stoll and colleagues (2020)

describe attentional selection of goal-relevant memories as one of the three essential pillars of their "retrieval mode": Attentional selection, along with intention to remember and the appropriate level of the neurotransmitter acetylcholine in the hippocampus, is described as necessary to be able to retrieve memories. According to this retrieval mode, attentional selection involves activity in posterior parietal cortex and ventrolateral prefrontal cortex.

The selectivity observed in the retrieval of memories bears similarity to the deployment of selective attention while searching for items in the external world. Chun and colleagues (2011) proposed that attention should broadly be categorised depending on the focus of attention to external vs. internal information. They propose that both external (perceptual) and internal (memory-oriented) attention share some core functional characteristics: the problem of limited capacity, a solution in form of selectivity, and the challenge of sustaining vigilance over temporal delays. Indeed, the encoding process of memory could be seen as the persisting traces of perceptual attention, whilst the retrieval process is activated by internal, reflective attention (Chun & Johnson, 2011). It could even be argued that the "ways in which we form, retrieve and work with memories largely represent acts of attention" (Kuhl & Chun, 2014). Would it be useful, therefore, to think of selective retrieval as a form of selective attention within the internal space of long-term memories?

Until fairly recently, selective attention was only investigated in relation to external stimulation. It was thought that other than for items within iconic memory (merely lingering within sensory stores, Averbach & Corriel, 1961; Neisser, 1967), all memories were impervious to selective attention. Even short-term, working memories were considered to be beyond interference or the influence of attention: Post-cues, presented at over 1 second post memory delay, were found ineffective at improving the

recall of items in visual short-term memory (Sperling, 1960). This assumption was challenged by the introduction of "retrocues": informative cues that retrospectively predicted the likely location of an item that would subsequently be relevant for memory-based performance conferred significant benefits (Griffin & Nobre, 2003; Landman et al, 2003). The retrocueing effect, providing evidence for an internal focus of attention, has been replicated manifold (review Souza & Oberauer, 2016; Nobre & Stokes, 2020). This benefit of orienting attention to memories has been observed across modalities (audition: Backer & Alain, 2012; Backer, Binns & Alain, 2015; tactile: Katus, Andersen & Müller, 2012a, 2012b, 2014), and even species (birds: Fongaro & Rose, 2020). The modulatory consequences of retrocues are measurable with neurophysiological measures (Kuo et al, 2012; Wallis et al, 2015) and visible within neuroimaging studies (Lepsien & Nobre, 2007; Sligte et al, 2009; Lepsien et al, 2011; Kuo et al, 2014; LaRocque et al, 2017).

Retrieval cues in long-term memory studies can also be considered as a type of retrocue: Comparably to retrocues, memory cues retrospectively point to specific aspects of a long-term memorandum. Recall of the cued item facilitates the retrieval of surrounding content until the whole memory episode or item is retrieved. This retrieval process, termed "pattern completion" (e.g., Rolls, 2013), is thought to constitute the mechanism by which the neural regions from encoding are re-activated (review Danker & Anderson, 2010). Is the act of this successful memory retrieval thus the consequence of orienting attention within LTM? It is notable that the recollection process in response to a cue could encompass both voluntary and involuntary elements. Moscovitch and colleagues (2016) describe the recollection process in response to a cue as containing two stages: an involuntary rapid stage (ecphory, thought to be dependent on the hippocampus) and a slower, conscious stage (dependent on cortical processes).

The popular Attention to Memory (AToM) theory specifically proposes a role for the internal orienting of attention during retrieval (Ciaramelli, Grady, & Moscovitch, 2008). An impetus for the development of the theory was the consistent finding of parietal cortex engagement during long-term-memory retrieval tasks. Whilst not traditionally considered a memory-related area, the parietal cortex was found to be activated during memory retrieval in EEG, PET, and fMRI studies (Rugg, 1995; Cabeza & Nyberg, 2000; for review see Cabeza et al 2008). The AToM theory proposed that the parietal cortex may be similarly important for controlling attention in memory as it is in controlling attention within the external environment. Paralleling the proposal that dorsal and ventral parietal cortex differentially contribute to goal directed and automatic attention respectively (Corbetta & Shulman, 2002), AToM proposes that these different subregions of the parietal cortex are involved in goal-based guidance of attention in memory and in the automatic capture of attention by spontaneous recollections. The dual attentional processes (DAP) hypothesis by Cabeza (2008) introduces a similar framework: Retrieval goals guide top-down attentional processes using the dorsal parietal cortex (DPC), whilst retrieval output captures bottom-up attention via the ventral parietal cortex (VPC).

However, to date, it is difficult to investigate possible contributions of selective attention to the retrieval of specific items or events within LTM, due to the lack of a rigorous experimental approach to isolate the systems and mechanisms specifically related to orienting attention in memory. Developing such an experimental approach will prove highly informative for addressing many issues regarding the nature of attentional selection in long-term memories.

Firstly, are the hallmarks of selective attention evident when focusing on internal items? In WM, attentional orienting leads to the selective prioritisation of the

relevant content during the maintenance period (Kuo et al, 2014; LaRocque et al, 2017; Lepsien & Nobre, 2007; Sligte et al, 2009; Kuo et al, 2012; Poch et al, 2014; Wallis et al, 2015). Would such prioritisation be observable for LTM? There is indirect evidence that selective attention modulates hippocampal activation during retrieval: Hashimoto and colleagues (2012) found that the novelty signal of the hippocampus was modulated depending on which object features participants attended to: Hippocampal activity was increased when objects semantically similar to studied ones were judged on a perceptual level (seen as novel), but not on a semantic level (seen as old). Further, Duncan and colleagues (2012) observed that the response of the hippocampal subfield CA1 as a match/mismatch detector was modulated depending on attention being directed to the relevant dimension. These studies support the idea that, in principle, selective attention can influence patterns of brain activity involved during retrieval.

A related question concerns the fate of the selected and ignored memoranda: WM studies have observed that attended memoranda are retrieved more quickly and accurately than unattended ones (standard retrocue RT effect, Griffin & Nobre, 2003). Tasks using continuous responses to reproduce perceptual qualities of memoranda have further shown that attentional selection can also improve the precision with which selected items are reported (e.g., Makovski & Pertzov, 2015; Gunseli et al, 2015; but see Souza & Oberauer, 2016). In comparison, unattended items have been observed to degrade in quality (LaRocque et al, 2015). Can focusing on an aspect of LTM similarly selectively improve the retrieval of relevant items and impair retrieval of unattended items? The most closely related evidence within long-term memory comes from studies using the Think/No Think paradigm, in which selected-against memory items are degraded (*suppression-induced forgetting*, Anderson & Hanslmayr, 2014).

Assuming that top-down attention-related mechanisms can guide successful retrieval, it becomes interesting to investigate what types or information are most useful for focusing internal search and what aspects of the mnemonic codes are most prone to modulation. Within the visual WM literature, proactive modulations in response to retrocues occur in multiple sensory areas and are also thought to modulate the ability of the memorandum to influence action, through output-gating processes (Myers et al, 2017, for discussion). Are any of the multiple levels of processing that may participate in memory retrieval – spanning from simple sensory features to complex relational information – particularly labile to the influence of selective internal attention?

In order to provide an experimental framework for investigating the psychological mechanisms of orienting attention to selective aspects of long-term memoranda, I developed a new experimental paradigm based on the retrocueing literature in working memory. The approach involves participants learning a set of novel associations between scenes and object identities and locations, and subsequently testing whether memory performance for selective objects within scenes can be specifically modulated by informative retrocues. Across three experiments, informative retrocues facilitated retrieval performance: Participants showed an improved recognition of the identity of objects associated with the scenes, as well as reproduction of their specific location. The experimental approach is highly adaptable and opens the doors to an exciting domain of questions regarding how attention guides selective retrieval.

Study 1: Developing an experimental approach to investigate selective attention during LTM retrieval

The development of this new experimental approach required an appropriate and well controlled learning task in which multiple items were associated with a given scene. Participants explored complex scenes, each of which contained two embedded novel items. A delayed forced-choice recognition and localisation task, probing the identity and location of each of the items, allowed participants to learn the identity and location of each object associated with a particular scene. Once participants had acquired the novel set of objects-in-scenes associations through repeated learning blocks, it was possible to test whether spatial retrocues during retrieval preferentially enhanced performance for the cued item within the scene. A separate, subsequent retrieval task evaluated the precision of the long-term memories and measured any influence of the preceding retrocueing task.

Method

The experimental methods were approved by the University of Oxford Central University Research Ethics Committee (CUREC), and the study was conducted in accordance with the Declaration of Helsinki.

Participants

In total, 22 healthy participants (13 female, 9 male; 20 right-handed; age range: 18 to 35; M = 23.9; SD = 5.4) took part in the study. All participants had normal or corrected-to-normal vision and reported no difficulties with colour vision. Participants gave written informed consent before testing and were financially reimbursed for their time at a rate of £10 per hour.

Materials

The main stimulus materials in the study comprised digital images of 138 scenes and 521 objects. The scene stimuli were collected by lab members or freely available in the public domain. Scenes were resized to 1000 x 750 pixels. The object images were collected from the public domain and drawn from a variety of object categories such as toys, utensils, food, and everyday objects. Objects were resized to 75 pixels in the horizontal or vertical dimension, depending on which was longest. A set of eight scenes and thirty objects was set aside for practice trials for all participants. Sixty-four scenes and 352 objects were randomly drawn from the remaining scenes and objects for the main tasks.

The experimental tasks were programmed and presented using Cogent toolbox (Cogent 2000, version 1.29 (RRID:SCR_015672)) running in Matlab (2015b, The MathWorks, Inc., Natick, Massachusetts, United States). The stimuli were presented using a screen resolution of 1920 x 1080 pixels (monitor: acer XB27OH 27 inch, screen width 60cm) and a 100-Hz refresh rate and were viewed from a distance of 90 cm. Scenes were presented in the centre of the screen on a grey background (RGB: 0.5 0.5 0.5). In different parts of trials and tasks, objects appeared either superimposed within scenes or against a grey background (RGB: 0.5 0.5 0.5).

Design

In the experiment, participants completed one demographic questionnaire and performed three tasks: learning, retro-cueing, and retrieval. The tasks are illustrated in *Figures 20 - 22 a*.

Learning task

In the Learning task (*Figure 20a*), participants searched through the scenes for two superimposed objects (one on each side of the scene). Their task was to find the objects within each scene and memorise their identities and locations, in order to report them back when probed later in the trial.

Each trial started with a study phase, in which a scene appeared for 5000 ms, and participants were free to explore them using eye movements. Object identities were selected randomly from the object set described above. Object locations were randomly drawn with the following restrictions: One object was presented on the right side of the scene while the other was presented on the left. Moreover, objects were presented at least 50 pixels away from edges of the scene and at least 100 pixels away from the vertical centre line of the scene. The scene presentation was followed by a blank delay (5000 ms).

Following each scene in the study phase, two probe phases were completed. In each of these forced-choice probes, two objects appeared, one 75 pixels to the left and one 75 pixels to the right of the fixation. Participants made a forced-choice response to select the target item that had been previously contained within the scene by clicking the corresponding left or right mouse button. The other object was a randomised, novel foil. The left vs. right positioning of the items was allocated randomly independently of whether the target object had previously appeared on the right or the left side of the scene. The choice stayed on the screen until a response was made. Following the choice, the scene reappeared on the screen without delay, and participants positioned the selected object within the scene, at its remembered location, by moving the mouse and confirming the choice with a press on the space bar. The second probe phase was

equivalent, but this time contained the other object from the scene and a different foil. The inter-trial interval, after completion of both probe phases, was 1000 ms.

The identities and locations of objects in scenes were studied over three experimental blocks, each containing all 64 scenes. The order in which scenes appeared in each block was randomised, as was the order in which the two scene-associated objects was probed. The object identities and positions within individual scenes were held constant across learning blocks for participants but differed across participants. The task design allowed for a within-subject analysis of the effect of learning block (3 levels) and order in which the objects were probed (2 levels). Dependent measures included identification accuracy, reaction time, and spatial error.

Retrocueing task

In the Retrocueing task (*Figure 21a*), participants viewed previously studied scenes without any embedded objects and, at the end of the trial, made a forced choice to select which of two objects in a probe array had previously been associated with the scene (in the Learning task). In between the presentation of the scene and of the probe array, an informative spatial retrocue indicated the side in which the relevant probe item had been learnt in the scene (100% validity) or an uninformative neutral retrocue provided no spatial information about the learned location of the relevant probe item.

Each previously learnt scene from the learning task was presented for 1000 ms without the objects, followed by a blank grey screen (750-ms duration). On half of the trials, this was followed by a spatially informative retrocue (spatial, 250-ms duration): One side of the fixation would turn blue (RGB: 0 0.5 1). This blue triangle pointed to the side of the screen in which the to-be-retrieved object had previously been learned. These retrocues were 100% informative (valid) of the upcoming probe. In the

remaining half of the trials, the retrocue was uninformative (neutral, 250-ms duration). In these trials, the fixation would turn completely blue, providing no information about the learned location of the item to be probed. Items previously learned on the right and left side of the scenes were equally likely to be probed for both retrocue conditions.

Following the retrocue, a blank grey screen appeared as a delay (2000-ms duration) before the presentation of a probe array containing two objects, with a distance of 75 pixels on either side of the fixation. One of the objects had previously been embedded in the scene during the Learning task (target) and the other object was a novel foil. Participants had to pick the object that had been associated with the scene during the Learning task by pressing the left or right mouse button corresponding to the target item on the screen. In trials with a valid spatial retrocue, the object learned at the cued location was always probed. In trials with a neutral retrocue, either learned object was equally likely to be probed. Target objects appeared randomly on the left or right of the fixation, independently of their original learnt position within the scene. Participants were instructed to respond as quickly and as accurately as possible. No feedback was given, and after an inter-trial interval of 1000 ms, the task moved on to the next trial. Only one of the two objects associated with each scene was probed at this stage of the study. The task design allowed a within-subject comparison of the effect of spatial versus neutral retrocues on the dependent measures of identification accuracy and reaction time.

Retrieval task

In the Retrieval task (*Figure 22a*), participants were prompted with each scene (with no embedded objects) and had to identify the associated objects and position them within the scene context. These forced-choice decisions were designed as a surprise

memory test to assess how well participants retained the learnt scene-object-location associations.

Each of the 64 scenes was presented for 1000 ms without any objects, followed by a 500-ms delay and then a probe array containing a previously scene-associated object and a randomised foil. The configuration of the probe array was similar to that used in the Learning and Retrocue tasks (75 pixels on either side of the fixation, with random location of learned object relative to original placement). Participants made a forced choice to select the probe object that had been learned in the scene by pressing the corresponding right or left mouse button. The scene reappeared, and participants placed the selected object within the scene by moving the mouse and confirming the choice by pressing the space bar. This process was repeated to test the memory for the other item learned within the scene. After both probe arrays, an inter-trial interval of 1000 ms followed. All learnt scenes and both associated objects were probed this way in a randomised order. Memory for half of the probed items had also been tested in the Retrocue task. In addition to providing a metric of overall memory quality, it was possible to assess any effects related to prior probing (2 levels: object probed in the Retrocueing task or not) and to retrocue type association (2 levels: scene associated with a spatial or a neutral retrocue). Dependent measures included identification accuracy, reaction time, and spatial error.

Procedure

Participants were informed that the study would consist of three computerised tasks and a computer-based questionnaire on demographic data. After reading the information sheet, participants provided informed consent. They were given time to read a task description including schematics of the Learning task before the

experimenter explained it again in person to allow for questions. Participants were informed that they had to memorise both object identity and location over the delay and were instructed to be as accurate and as fast as possible. Participants completed a block of practice trials before moving on to the three blocks of the Learning task. After each block, participant could take a short break. When ready, they initiated the next block by pressing the space bar. The Learning task took approximately one hour to complete.

After the Learning task, participants had a 20-minute break. They completed a demographic questionnaire and otherwise were free to spend their break however they liked. Participants were unaware that they would be required to recall the learnt information in the tasks following the break.

After the break, participants were given a written and illustrated description and set of instructions for the Retrocueing task to study in their own time. Subsequently, the experimenter went through the task with them, allowing for questions. The instructions and the experimenter stressed the importance of responding as quickly and as accurately as possible. Participants then completed a practice task supervised by the experimenter. If participants performed the practice task accurately, they moved on to the Retrocueing task, which took approximately 10-15 minutes to complete. Should participants be unable to comprehend the task demands, their data would be excluded from the analysis later. All participants were able to complete the practice task satisfactorily.

Once they had completed the Retrocueing task, participants were given written instructions and a schematic of the Retrieval task to read. The experimenter went through the task again with the participants, allowing for questions. As the task demands were very similar to the Learning task, no computerised practice task was

required. Participants completed the Retrieval task, which lasted approximately 10-15 minutes.

Once they completed all tasks, participants had the opportunity to be debriefed about the study aims. Lastly, they completed a payment form. In total the study took just under two hours to complete.

Analysis

Learning task

The dependent measures of interest in the Learning task were object identification accuracy and response times (RT), as well as the spatial error measured as Euclidean distance (in pixels) between the response location and the correct location of the target object within the scene. These measures were calculated per object – first choice, second choice – and for each of the three blocks. For statistical analyses, accuracy data were transformed using arcsine, and RT and spatial-accuracy data using log. I conducted separate repeated-measures ANOVAs for RT and spatial accuracy data to test for effects of block (3 levels: 1^{st}, 2^{nd}, 3^{rd}) and object (2 levels: probed 1^{st} or 2^{nd}). Trials were excluded based on inaccurate target identification or reaction times (RTs) slower than 3 standard deviations of the condition mean.

For spatial-error data, trials were excluded if participants placed the object outside of image bounds. Chance performance for the spatial-error data was estimated for each participant, using a permutation approach. Chance performance was estimated by calculating the average distance of hypothetical object-response pairs, by shuffling x and y coordinates of actual object locations and responses over 1000 iterations. The averaged distance of all potential object-location response pairs gave an estimate of a

chance level of spatial error: It mimicked the spatial error if responses were given randomly and irrespective of the target coordinates. This measure was estimated in each of the Learning, Retrocueing and Retrieval task for each participant.

Retrocueing task

In the Retrocueing task, the dependent measures of interest were object identification accuracy and reaction times. For statistical analyses, accuracy data were transformed using arcsine, and RT data using log. Paired t-tests tested for the effect or spatial vs. neutral retrocues. Further, I analysed whether the hemispherical congruency between learnt target location and target appearance during the Retrocueing task affected behavioural measures, and whether this interacted with the effect of the retrocue. This was assessed in a 2 (type of retrocue) by 2 (congruent or incongruent) repeated measures ANOVA, or equivalent, if assumptions of normality were not met.

Data were excluded for scenes where the wrong object had been chosen in block 3 of the Learning task, since the object-scene association was considered as not having been learnt. On average, 1.70% (*SE* = .38%) of trials were excluded this way per person. For the RT analysis only, trials were also excluded based on identification accuracy or if reaction times were beyond 3 standard deviations away from the mean within participants in that conditions.

Retrieval task

The dependent measures of interest in the Retrieval task were object identification accuracy and reaction times, as well as the spatial error in object localization. The experimental factors were based on the Retrocueing task: if the scene had appeared with a spatial or neutral retrocue (retrocue condition) and if the object had been probed or

unprobed (probe condition). These conditions were used as factors in a 2 x 2 repeated-measures ANOVA. Prior to statistical analyses, the data were transformed using arcsine (accuracy) and log (RT, spatial error). For the accuracy analysis, trials were excluded if participants had not found the object in the third block of the learning task. For the RT and spatial-error analysis, only trials with correct object identification were included. Trials were also removed if response times were exceptionally fast or slow (more than ± 3 SD of the mean) within the Recall task. Lastly, spatial error data were excluded if the object was placed outside the image bounds.

Results

Learning task

Behavioural performance in the Learning task is summarised in *Figure 20*. Participants were highly accurate in choosing the correct objects studied within the scenes, and accuracy further increased systematically across the three blocks. Reaction time and spatial error decreased systematically over the three blocks, indicating good learning of object locations within scenes.

A 3 x 2 repeated-measures ANOVA on accuracy data showed a significant main effect (F (2, 42) = 9.72, $p < .001$, $\eta_p^2 = .32$) and linear contrast (F (1, 21) = 18.10, $p < .001$, $\eta_p^2 = .46$) of block, indicating that accuracy increased successively over the three learning blocks (M, SE (in %): Block 1 = 96.73, .48; Block 2 = 97.94, .38; Block 3 = 98.62, .29). There was no main effect of object (F (1, 21) = .06, $p = .808$, $\eta_p^2 < .01$), and no interaction between block and the order of probed object (F (2, 42) = .74, $p = .481$, $\eta_p^2 = .03$), *Figure 20b*.

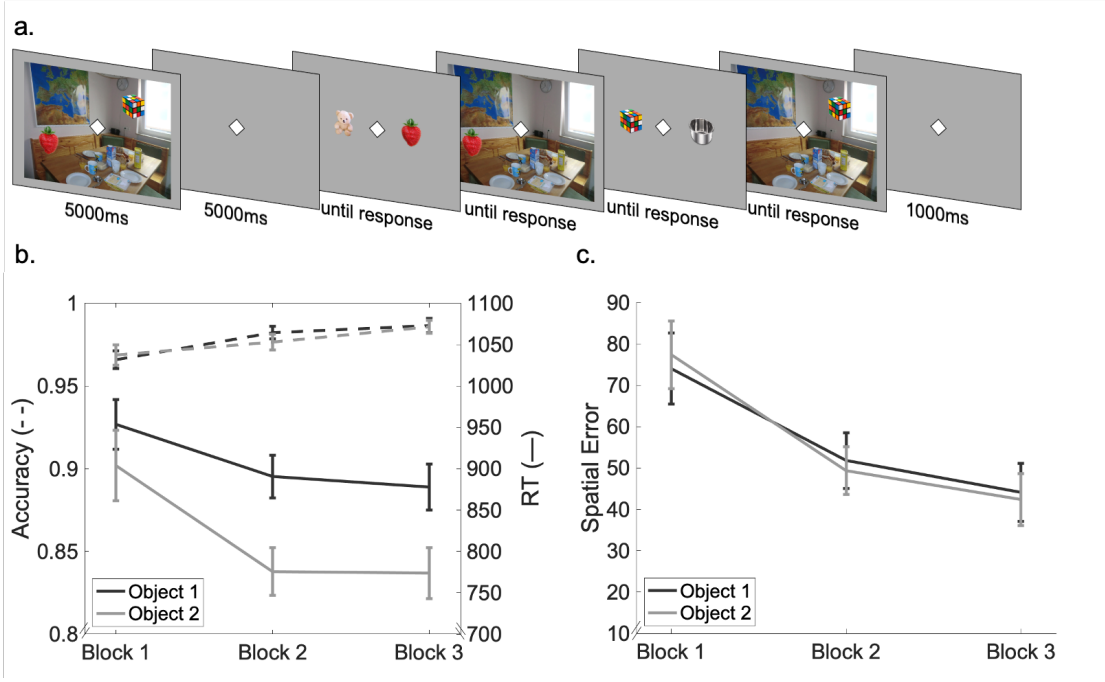


Figure 20. Panel a. illustrates the Learning task. Panel b. shows the increasing identification accuracy (dotted lines, in proportion of accurately chosen objects) and decreasing reaction times (continuous lines, in ms) over the three learning blocks. Dark grey depicts measures relative to the first object probed and light grey to the second object. Error bars are always standard error of the means. Panel c. shows the decreasing spatial error in pixels when positioning the two objects in the three learning blocks (chance level: $M = 427.49$, $SE = 2.61$).

A 3 x 2 repeated-measures ANOVA on the reaction time data resulted in significant main effects of block (F (2, 42) = 25.92, $p < .001$, $\eta_p^2 = .55$) and object (F (1, 21) = 15.70, $p < .001$, $\eta_p^2 = .43$). Overall, RTs for the second object that was probed ($M = 817.75$, $SE = 32.90$) were faster than for the first one ($M = 907.38$, $SE = 24.46$), *Figure 20b*. Both objects showed an overall decrease in RTs over the three blocks, as indicated by a significant linear contrast of block: F (1, 21) = 35.26, $p < .001$, $\eta_p^2 = .63$ (M, SE (in ms): Block 1 = 928.75, 32.63; Block 2 = 833.10, 24.81; Block 3 = 825.85, 26.62). Additionally, there was a significant interaction (F (1.51, 31.60) = 5.42, p = .015, $\eta_p^2 = .21$), suggesting that the difference in reaction times for the first and second object differed across the three learning blocks (*M(Object 1) – M(Object 2):* Block 1 =

49.87; Block 2 = 115.07; Block 3 = 103.97). Post-hoc tests confirmed that differences in RT were more accentuated for blocks 2 and 3 than for block 1. RTs in response to the two objects did not differ significantly for block 1, t (21) = 1.83, p = .081, d = .39, but they did for block 2 and 3, t (21) = 4.99, p < .001, d = 1.06 and t (21) = 4.34, p < .001, d = .93, respectively.

Participants' positioning of the objects within context was accurate overall, with the averaged spatial error in placement (M = 56.53 pixels, SE = 6.79) over all blocks being much lower than chance (M = 427.49 pixels, SE = 2.61), t (21) = -19.59, p < .001, d = 4.18. The 3 x 2 repeated-measures ANOVA showed a significant main effect (F (2, 42) = 51.37, p < .001, η_p^2 = .71) and linear contrast (F (1, 21) = 73.20, p < .001, η_p^2 = .78) of block, showing that spatial error decreased over the three learning blocks (M, SE (in pixel distance): Block 1 = 75.73, 8.29; Block 2 = 50.60, 6.18; Block 3 = 43.26, 6.58), *Figure 20c*. There was no main effect of object (F (1, 21) = .09, p = .762, η_p^2 < .01) and no interaction between object and block (F (1.56, 32.83) = 1.30, p = .280, η_p^2 = .06).

Retrocueing task

Analysis of the Retrocueing task showed excellent accuracy to select objects previously associated with a scene and a significant benefit conferred by spatial retrocues for response times to select the relevant item. A paired-samples t-test showed no difference in identification accuracy between spatial and neutral retrocue trials, t (21) = .25, p = .802, d = .05 (M, SE (in %), spatial: 95.58, .98; neutral: 95.65, 1.14), *Figure 21b*. However, for reaction times, a paired-samples t-test showed that there was a significant difference between conditions, t (21) = 4.95, p < .001, d = 1.06, *Figure 21c*: The reaction times (in ms) to select items were significantly shorter in spatial

retrocue trials (M = 735.91, SE = 44.05) than in neutral retrocue trials (M = 819.71, SE = 48.92).

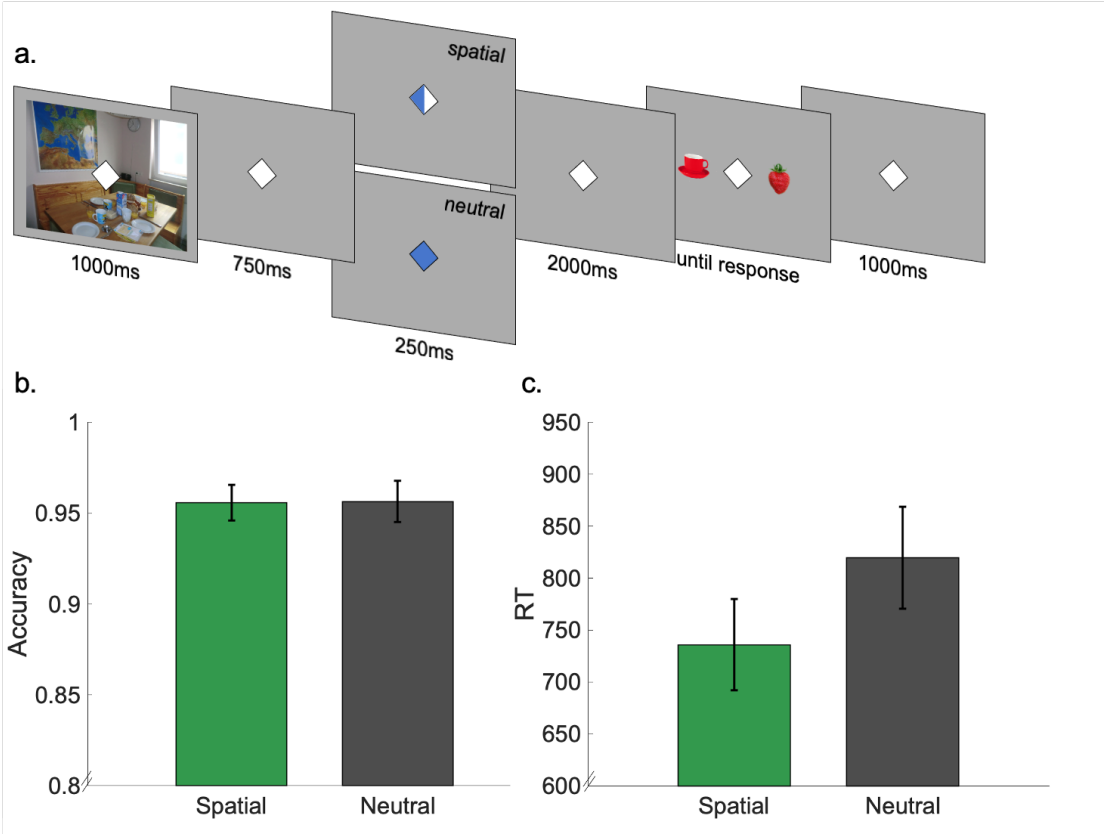


Figure 21. Panel a. illustrates the Retrocueing task. Panel b. shows the identification accuracy (in proportion of accurately chosen objects) for the spatial- (green) and neutral- (grey) retrocue conditions. No statistically significant difference in accuracy was observed. Error bars represent standard error of the means. Panel c. shows the significant decrease in reaction times (in ms) for spatially informatively retrocued compared to neutrally retrocued objects ($p < .001$).

Measures of congruency showed an interesting pattern highlighting the effects of congruency on reaction times for neutral but not spatial retrocues. In addition to the expected main effect of retrocue type, F (1, 21) = 39.55, $p < .001$, $\eta_p^2 = .65$, there was also a main effect of congruency, F (1, 21) = 8.49, $p = .008$, $\eta_p^2 = .29$, showing longer RTs for incongruently appearing targets (M = 806.83, SE = 56.60) in contrast to congruent target locations (M = 746.04, SE = 40.13). Importantly, there was also a significant interaction, F (1, 21) = 10.02, $p = .005$, $\eta_p^2 = .32$. Post hoc t-tests confirmed

that this interaction was driven by the absence of significant congruency effects following spatial retrocues, t (21) = -1.03, p = .314, d = .22 (congruent: M =713.44, SE = 42.22, incongruent: M = 732.02, SE = 52.41). Neutral retrocues were followed by significantly longer RTs for incongruent (M = 881.64, SE = 64.69) as opposed to congruent locations (M = 778.64, SE = 40.04), t (21) = -3.30, p = .003, d = .70. There was no effect of congruency observed for accuracy, showing no significant differences between congruent and incongruent locations in response to spatial or neutral cues, Z = -1.18, p = .237 and Z = -.38, p = .701, respectively.

Retrieval task

Performance in the Retrieval task was highly accurate and consistent across conditions with a mean of 97.48% (SE = .58%). A 2 x 2 repeated-measures ANOVA on accuracy data showed no main effect of probe (F (1, 21) = .66, p = .425, η_p^2 = .03) or retrocue condition (F (1, 21) = .72, p = .405, η_p^2 = .03), nor an interaction (F (1, 21) = .57, p = .461, η_p^2 = .03), *Figure 22b*.

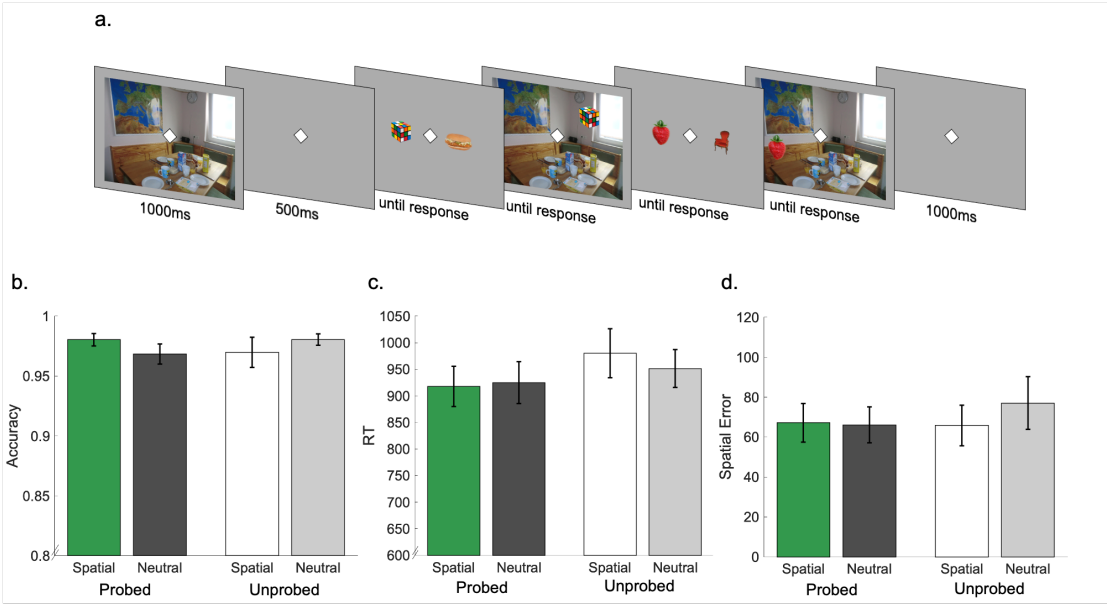


Figure 22. Panel a. depicts the Retrieval task. Panel b. shows that there was no difference in accuracy, irrespective of the two conditions "probe" (object was probed in the Retrocueing task or not) and "retrocue type" (scene was associated with a spatially informative [green/white] or neutral retrocue [dark/light grey]). Panel c. illustrates faster reaction times to probed objects in contrast to unprobed objects (p = .026). Panel d. shows no difference in spatial accuracy for positioning objects within the scene by conditions. Error bars represent standard error of the means.

Analysis of reaction times showed an advantage for objects that were previously probed during the retrocue task. A 2 x 2 repeated-measures ANOVA on reaction time data showed a significant effect of probe (F (1, 21) = 5.75, p = .026, η_p^2 = .22): Probed items were recalled faster (M = 921.21 ms, SE = 38.00) than unprobed items (M = 965.95 ms, SE = 38.65), *Figure 22c*. There was no significant effect of retrocue (F (1, 21) = .31, p = .584, η_p^2 = .01) and no interaction (F (1, 21) = .74, p = .400, η_p^2 = .03).

Spatial errors when positioning previously studied items within scenes showed that spatial memory was much better than chance (M = 428.43 pixels, SE = 2.60): Pairwise t-tests showed that the spatial errors for probed (spatial retrocue: t (21) = -13.52, p < .001, d = 2.88, neutral retrocue: t (21) = -14.11, p < .001, d = 3.01) and unprobed objects (spatial retrocue: t (21) = -14.39, p < .001, d = 3.07, neutral retrocue: t (21) = -11.38, p < .001, d = 2.43) were all significantly smaller than estimated chance values. The 2 x 2 repeated-measures ANOVA on spatial error yielded no significant effects (Probe: F (1, 21) = .07, p = .797, η_p^2 < .01, Retrocue: F (1, 21) = .43, p = .518, η_p^2 = .02, Interaction: F (1, 21) = .27, p = .608, η_p^2 = .01), *Figure 22d*.

Discussion

This first study provides a promising foundation for developing a rigorous experimental approach for investigating the ability to orient attention in LTM. Participants were able to learn a sizeable set of novel associations of multiple object identities and locations within scenes. Despite limited exposure times, participants

retained an accurate representation of these associations over a significant delay. Using spatial retrocues, it was possible to measure benefits of orienting attention to the location of a memorandum in LTM.

The Retrocueing task, emphasising speeded recognition of remembered items, thus provides a good basis for investigating benefits (and potentially costs) of orienting attention in memory. Within this study, the RT measures were the most sensitive in the Retrocueing task: Speeded responding was repeatedly stressed, and the memory discrimination was a simple forced-choice task between two dissimilar objects, one being an unrelated, randomised foil. The targets were well learnt and easily discriminable from unrelated foils. All stimuli were clearly visible and presented until response. There was a clear benefit in identification speed if a spatially informative retrocue oriented attention to the location, and hence the identity, of to be remembered memorandum. In order to use the retrocue efficaciously, participants had to remember which object had been learnt on the retrocued side of the scene, hence the learnt location of the object. Further, having an informative spatial retrocue eliminated the effects of congruency on RTs, emphasising the attentional benefit of the retrocue for using the selected LTM content. To our knowledge, this is the first study demonstrating that retrocues can be effective in long-term memory.

I observed no accuracy benefits in identifying the target identity after an informative retrocue. However, performance was very high, so the current measure may have not been sensitive enough to pick up on any potential effects because of possible ceiling-level performance. The sensitivity of the Retrocueing Task could therefore be adjusted by implementing more challenging retrieval conditions, allowing for larger scope in which potential accuracy effects could be expressed. Future experiments could utilise a brief exposure of the probe array, a presence vs absence judgment of the target

within a more complex array (see Patai et al, 2012) or indeed a precision measure in form of a spatial placement or a continuous delayed-estimation task (Richter et al, 2016). Rather than adjusting the Retrocueing Task, difficulty could also be increased by altering the Learning Task: If the association are learnt over fewer repetitions and thus less readily available in memory, it might increase the value of utilising the retrocues.

In this study, retrocues facilitated speeded responses and hence access to the memorandum in question. Yet, memory precision was not assessed. Changes in the task design such as the use of continuous performance variables (such as spatial placement, or colour matching) and increasing levels of task demands could give further insight into the nature of the retrocue benefit. Higher task demands might tease apart if the beneficial impact of retrocues stems only from a privileged access or if it may potentially improve the quality of the memorandum itself. Location precision, for example, could be probed by a spatial accuracy judgment of the target appearing in scene context at varying degrees of spatial error, or by a spatial placement task. Whilst it might be more sensitive to replace the identification task with a precision task, adding a precision task to the existing Retrocueing Task layout has the benefit of retaining the ability to detect improvements in the speed of identification (access) while also probing for improvements in spatial placement (precision).

Our Learning task created a useful set of memoranda with which to explore attention during retrieval. By associating two novel objects with each scene, I set up competing spatial/identity-object associations within each contextual memory and was able to manipulate attention to one or the other item. Performance in both the identification and the localisation of the objects was great and learning across blocks was efficient. Merely three exposures per objects-scene association were sufficient to

create a stable memory. Participants reported the task to be simple and engaging, which is promising for use in patient populations or developmental studies. Within this study spatial retrocues were used. In theory, it would also be possible to use retrocues based on object categories or other object-related features if these were manipulated properly during learning. It would be very simple to accommodate this manipulation within the present task set-up.

The accuracy measures in the Retrieval task ensured that participants were actually able to retain the learnt objects-locations associations over a prolonged time, and opened avenues of evaluating the mediating effects the retrocues might have on the stability of the acquired LTMs. One may have expected protracted consequences of directing attention to one item during the Retrocueing task on performance during the subsequent retrieval task. For example, Wimber and colleagues (2015) proposed that selecting one item in an episode could lead to weakening other competing ones. In our task, such a mechanism would have led to significantly poorer retrieval for unprobed items when spatial retrocues had oriented attention to the competing item during previous the retrocue task compared to when neutral retrocues had been used. Although a hint of this pattern was observed numerically, there was no significant statistical effect. Instead, the only observed effect was speeded access to the retrieved item in the previous task irrespective of retrocue. This is in line with the well-known testing effect (Abbot, 1909; Carrier & Pashler, 1992): Retrieval facilitates learning and improves memory. If potential effects of retrocues on retrieval are smaller and harder to evaluate, it might be worth increasing the sensitivity of the measures or the sample size for better power to get at these effects.

There are many important ways in which these initial findings remain to be validated in further studies. The timing of the trials in the Retrocueing task was very

slow, which may have de-emphasised the utility of using selective attention to guide retrieval. With the current absence of neural measures, it is difficult to prove that the RT effect was indeed due to proactive prioritisation. There may be an argument that the spatially informative retrocue might facilitate retrieval of the target item through a process of matching and verification according to compatibility with the cued location (cf. priming literature, e.g., Neely, 1991). The benefit of the retrocue would thus not be due to influence prior to target presentation, but after the probe items are presented and a choice has to be made between target and foil. Tracking modulatory effects using neural measures could tease apart whether anticipatory biases towards the retrocued target would develop proactively prior to the presentation of the probe array, or whether the difference appears only post-target presentation. Without neural measures, the robustness of the Retrocueing benefit in LTM could be tested by increasing the behavioural difficulty of the task: Speeding up the task, and in particular the exposure durations, would be the next important step. Study 2 addresses this by using faster intervals and briefer exposure durations.

Assuming that the retrocues can still be used efficaciously to guide retrieval under speeded conditions, the issue of access versus precision remains: Does attentional selection during retrieval only facilitate the access to the memory or can it influence the quality of the memory (outlined above)? Study 2 and 3 address this question by including a spatial placement task as a precision measure after the object choice in the Retrocueing task.

The comparison of trials with spatial versus neutral retrocues suggested benefits of spatial selective attention operating within LTM. However, it is not possible to rule out other nuisance effects. Neutral and spatially informative retrocues might inadvertently cause a different overall state, potentially with relation to arousal,

anticipated task difficulty, or processing of informative content. It is conceivable that a difference in general state might influence and confound the results attributed to the nature of the retrocue. Experiment 3 addresses this confound by testing recognition in trials with only informative cues, comparing performance when the target was the retrocued item (valid retrocue) or the other item learned in the scene (invalid retrocue).

Study 2: Replicating benefits of selective attention in LTM retrieval

In Study 1, a promising experimental approach was developed to investigate the effects of selective attention on the retrieval of long-term memories. Spatial retrocues allowed speeded access to the to-be-retrieved object, thus extending the retrocueing effect known in the WM literature to long-term memories. This current Study 2 aimed to replicate and extend the findings of Study 1 using an appropriate sample size based on power-calculations using the effect sizes from Study 1 and by increasing the performance demands of the retrocueing task.

Within Study 2, the Retrocueing task was adapted in two ways to address some of the issues discussed above. Firstly, I implemented speeded task conditions, by decreasing exposure durations and delays. Operating more closely at the perceptual and cognitive limits should test the reliability of the behavioural benefit of the retrocue. The faster pace was aimed at encouraging participants to utilise the retrocue to anticipate the probed object and proactively to prepare for the forced-choice decision. Secondly, I introduced a localisation response following object identification in the retrocue task. The spatial placement acted as a precision measure to provide information on whether the retrocues influenced memory precision as well as access to the memory.

Method

The experimental methods were approved by the University of Oxford Central University Research Ethics Committee (CUREC), and the study was conducted in accordance with the Declaration of Helsinki.

Participants

In total 40 healthy participants took part in the study (23 female, 15 male, 2 other; 34 right-handed; age range: 18 to 33; $M = 24.9$; $SD = 4.3$). All participants had normal or corrected-to-normal vision and reported no difficulties with colour vision. They gave written informed consent before testing and were financially reimbursed for their time at a rate of £10 per hour.

Materials

The stimuli were the same as those used in Study 1. All stimuli were presented using Cogent 2000 (RRID:SCR_015672, version 1.29), which ran via Matlab (2018, The MathWorks, Inc., Natick, Massachusetts, United States). Study 2 was run on a 23-inch screen (screen width 50.8 cm, resolution 1920 x 1080 pixels), with a 60 Hz refresh rate. The viewing distance was variable as participants were free to move but averaged around 60cm.

Design

Schematics of tasks procedures are depicted in *Figure 23 - 25 a*. The Learning and Retrieval tasks were the same as in Study 1. The Retrocueing task was adapted to increase its tempo and to increase the sensitivity of the memory measures. In the Retrocueing task (*Figure 24a*), scenes were presented for only 100 ms followed by a delay of 750 ms. The duration of the retrocue was shortened to 100 ms and the delay preceding the probe stage to 750 ms. Furthermore, following the identification of the target object in the Retrocueing task, the scene reappeared, and participants had to indicate the remembered location of the object on the scene. For this, participants had

to move the mouse, indicated by a white dot, to the remembered location, and then to confirm the selection by pressing the space bar.

Participants completed two sets of practice trials before the main Retrocueing task. The first set of practice trials used the practice stimuli to familiarise participants with the sequence of events and required responses. The second set of practice trials used two scenes from the Learning task and their associated objects. In these practice trials, each of the scene-objects associations were used four times to practise the task with spatial and neutral retrocues. The practice allowed participants to familiarise themselves with having to recall objects quickly, choosing the remembered objects fast and correctly, and then positioning them accurately in the scene context.

Procedure

The experimental protocol was the same as in Study 1 except for the following changes: Considering the increased difficulty of the Retrocueing task, due to its increased pace, participants completed two sets of practice trials prior to the Retrocueing task (as described above). During all tasks, participants were encouraged to respond as accurately and as quickly as possible.

Analysis

The analyses of the data from the Learning and Retrieval tasks were the same as described in Study 1. The majority of the data analysis of the Retrocueing task was also equivalent. In addition, the new Retrocueing task included a measure of spatial error measured as Euclidean distance (in pixels). The same analysis procedures were used as for the equivalent measure in the Retrieval task.

Results

Learning task

Behavioural performance in the Learning task is summarised in *Figure 23*. The results from the Learning task replicated the pattern observed in Study 1: Participants were very accurate in choosing the objects studied within the scenes, and accuracy further increased progressively across the three blocks. Reaction time and spatial error decreased systematically over the three blocks, indicating good learning of object locations within scenes.

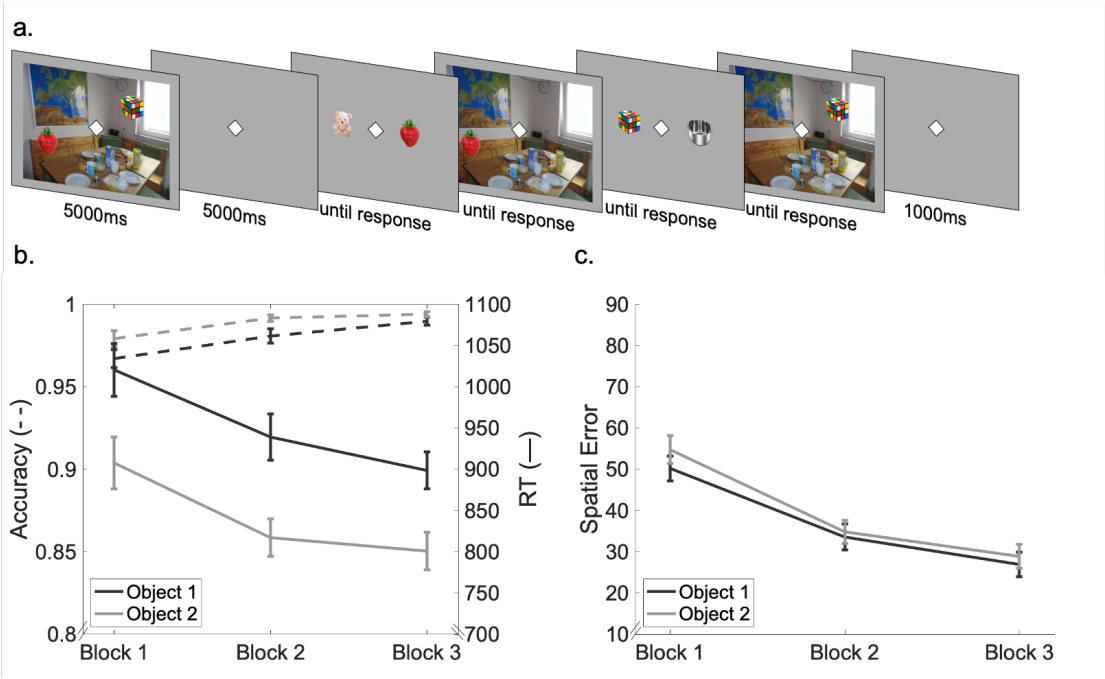


Figure 23. Panel a. depicts the Learning task. Panel b. shows the increasing identification accuracy (dotted lines, in proportion of accurately chosen objects) and decreasing reaction times (continuous lines, in ms) of participants over the three learning blocks. Dark grey illustrates measures relative to the first probed object, and light grey to the second object. Error bars are standard error of the means. Panel c. shows the decreasing spatial error in pixels when positioning the two objects in the three learning blocks (chance level: $M = 428.51$, $SE = 1.94$).

Comparable to Study 1, a 3 x 2 repeated-measures ANOVA on the accuracy data demonstrated a main effect ($F\,(2, 78) = 19.48$, $p < .001$, $\eta_p^2 = .33$) and a linear contrast ($F\,(1, 39) = 34.89$, $p < .001$, $\eta_p^2 = .47$) of block. These indicated that accuracy increased

systematically over the three blocks of learning (*M*, *SE* (in %): Block 1 = 97.32, .44; Block 2 = 98.63, .29; Block 3 = 99.20, .17), *Figure 23b*. There was also a main effect of object (F (1, 39) = 18.62, $p < .001$, $\eta_p^2 = .32$). The second probed object was chosen with greater accuracy ($M = 98.84$, $SE = .25$) than the first object ($M = 97.93$, $SE = .31$). There was no interaction between blocks and objects (F (2, 78) = .89, $p = .417$, $\eta_p^2 = .02$).

A 3 x 2 repeated-measures ANOVA on reaction times resulted in significant main effects of block (F (1.53, 59.63) = 33.06, $p < .001$, $\eta_p^2 = .46$) and object (F (1, 39) = 30.90, $p < .001$, $\eta_p^2 = .44$), *Figure 23b*. Reaction times decreased systematically over blocks, as indicated by a significant linear contrast: F (1, 39) = 40.55, $p < .001$, $\eta_p^2 = .51$ (*M*, *SE* (in ms): Block 1 = 964.33, 29.54; Block 2 = 878.18, 22.74; Block 3 = 849.79, 20.20). Reaction times were generally faster for the second object ($M = 841.98$, $SE = 23.92$) in comparison with the first object ($M = 952.89$, $SE = 25.57$), *Figure 23b*. There was no significant interaction (F (2, 78) = .61, $p = .546$, $\eta_p^2 = .02$).

Participants were able to position the objects accurately within scenes. The overall averaged spatial error ($M = 38.22$ pixels, $SE = 2.74$) was much lower than chance ($M = 428.51$ pixels, $SE = 1.94$), t (39) = -35.30, $p < .001$, $d = 5.58$. The 3 x 2 repeated-measures ANOVA on spatial error showed a significant main effect (F (1.62, 62.98) = 79.14, $p < .001$, $\eta_p^2 = .67$) and linear contrast (F (1, 39) = 104.04, $p < .001$, $\eta_p^2 = .73$) of block, indicating that spatial error decreased systematically over the three learning blocks (*M*, *SE* (in pixels): Block 1 = 52.52, 3.14; Block 2 = 34.22, 2.91; Block 3 = 27.91, 2.90), *Figure 23c*. There was also a main effect of object (F (1, 39) = 16.12, $p < .001$, $\eta_p^2 = .29$), such that the first objects were associated with a lower spatial error ($M = 36.93$, $SE = 2.74$) in contrast to the second objects ($M = 39.51$, $SE = 2.79$). There was no significant interaction (F (2, 78) = .59, $p = .558$, $\eta_p^2 = .02$).

Retrocueing task

Analysis of the Retrocueing task showed high accuracy to select and position objects previously associated with a scene as well as a significant benefit conferred by spatial retrocues for response times to select the relevant item. The spatially informative retrocue was again associated with decreased congruency effects.

A paired samples t-test comparing identification accuracy showed no difference by retrocue condition: t (39) = -1.50, p = .142, d = .24, *Figure 24b*. In both cases item identification was highly accurate (*M*, *SE* (in %), spatial: 92.29, 1.56; neutral: 90.88, 1.37).

Reaction-time data provided sensitive measures revealing beneficial effects of retrocueing. A paired-samples t-test resulted in a significant difference by retrocue condition, t (39) = 2.93, p = .006, d = .46: Participants were significantly faster to recognise memory targets following a spatially informative retrocue (*M* = 748.88 ms, *SE* = 30.83) than a neutral retrocue (*M* = 775.81 ms, *SE* = 25.80), *Figure 24c*.

The spatial error for positioning the probed objects in the scene was much smaller than chance (*M* = 431.59 pixels, *SE* = 2.15) after both spatially informative (*M* = 63.41 pixels, *SE* = 7.88) and neutral retrocues (*M* = 73.59 pixels, *SE* = 9.14), t (39) = -21.74, $p < .001$, d = 3.44 and t (39) = -17.44, $p < .001$, d = 2.76, respectively. Although spatial errors were numerically lower in trials with spatial retrocues, a paired samples t-test exhibited no difference in spatial error between retrocue conditions: t (39) = 1.42, p = .163, d = .22, *Figure 24d*.

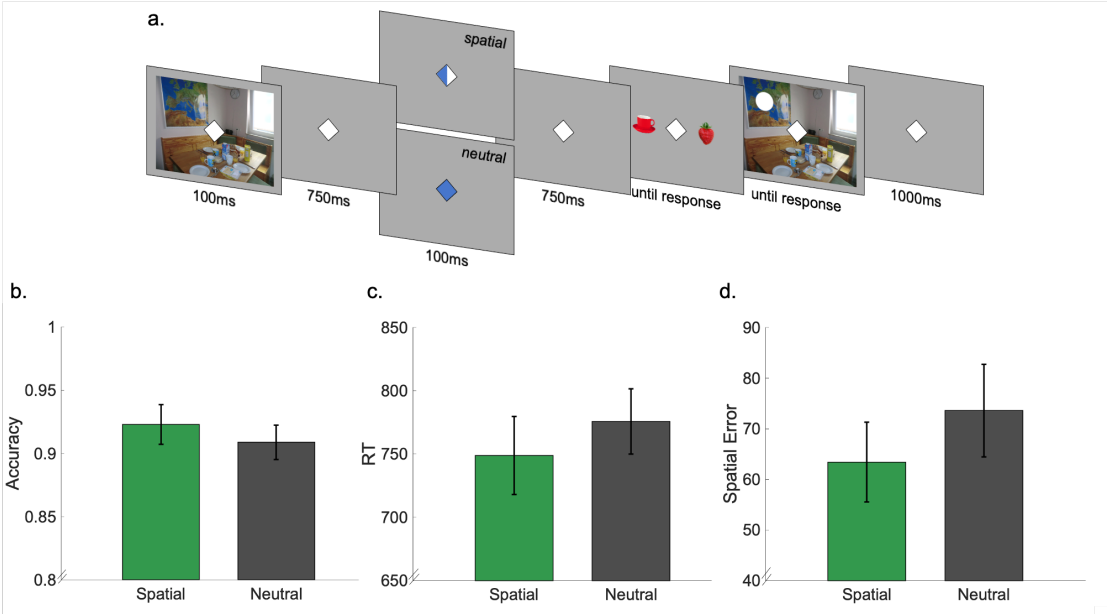


Figure 24. Panel a. illustrates the Retrocueing task. Panel b. shows the identification accuracy (in proportion of accurately chosen objects) for the spatial- (green) and neutral- (grey) retrocue conditions. No statistically significant difference in accuracy was observed. Panel c. depicts the decrease in reaction time (in ms) to spatially informatively retrocued compared to neutrally retrocued objects ($p = .006$). Panel d. shows the level of spatial error (in pixels) when indicating the remembered location of the chosen object. There was no significant difference between retrocueing conditions. Error bars all represent standard error of the means.

Breaking these effects down by the spatial congruency of the learnt and testing location, the pattern was more nuanced, emphasising the benefit of the spatial retrocue. With respect to identification accuracy, incongruent target locations were only associated with a worse identification rate if following neutral retrocues, Z = -2.58, p = .010 (congruent: $M = 92.69$, $SE = 1.49$, incongruent: $M = 88.80$, $SE = 1.67$), as opposed to spatial retrocues, Z = -.08, p = .938 (congruent: $M = 94.05$, $SE = 1.48$, incongruent: $M = 93.76$, $SE = 1.49$). Indeed, testing for a retrocueing benefit depending on congruency, the spatial retrocue was associated with more accurate identification in incongruent trials only, $Z = -3.07$, $p = .002$. Such a retrocueing benefit was not observed in congruent trials, $Z = -.90$, $p = .368$.

The results of congruency on reaction time replicated the observed results in Study 1: There was the expected main effect of retrocue type, F (1, 39) = 19.44, p <

.001, η_p^2 = .33, as well as a main effect of congruency, F (1, 39) = 11.93, p = .001, η_p^2 = .23, showing longer RTs for incongruent (M = 762.37, SE = 28.23) than congruent target locations (M = 735.21, SE = 25.22). Importantly, as a significant interaction, F (1, 39) = 5.76, p = .021, η_p^2 = .13, indicated, the congruency effects were subject to modulation depending on the retrocue. The congruency effect was only significant following neutral retrocues, t (39) = -3.00, p = .005, d = .47, whilst the equivalent contrast following spatial retrocues was not significant, t (39) = -1.85, p = .072, d = .29.

With respect to spatial error, there was no main effect of cue type, F (1, 39) = 2.57, p = .117, η_p^2 = .06, but again a main effect of congruency, F (1, 39) = 4.35, p = .044, η_p^2 = .10, with larger errors for incongruent (M = 72.90, SE = 9.23) than for congruent target locations (M = 63.44, SE = 7.55). The effects of congruency differed depending on retrocue association, indicated by a significant interaction, F (1, 39) = 4.60, p = .038, η_p^2 = .10. Post hoc t-tests confirmed that there was a significant congruency effect following a neutral retrocue, t (39) = -2.16, p = .037, d = .34 (congruent: M = 64.48, SE = 7.85, incongruent: M = 83.77, SE = 11.97), yet this effect was absent following a spatially informative retrocue, t (39) = .42, p = .677, d = .07 (congruent: M = 62.39, SE = 8.07, incongruent: M = 62.04, SE = 7.86). Further, when taking congruency into account, there was a significant retrocueing effect on spatial error in incongruent trials, t (39) = -2.22, p = .032, d = .35, with smaller errors following a spatial retrocue. This pattern was not observed for congruent trials, t (39) = -.39, p = .702, d = .06.

Retrieval task

Performance in the Retrieval task was highly accurate and consistent across conditions with a mean of 98.31% (SE = .48%). A 2 x 2 repeated-measures ANOVA

on accuracy data showed no main effect of probe (F (1, 39) = .19, p = .663, η_p^2 < .01) or retrocue condition (F (1, 39) = .02, p = .895, η_p^2 < .01), and no interaction (F (1, 39) < .01, p = .949, η_p^2 < .01), *Figure 25b*. There were also no significant effects in reaction time data, analysed in a 2 x 2 repeated-measures ANOVA (Probe: F (1, 39) = 1.23, p = .274, η_p^2 = .03, Retrocue: F (1, 39) = .10, p = .749, η_p^2 < .01, Interaction: F (1, 39) < .01, p = .989, η_p^2 < .01), *Figure 25c*.

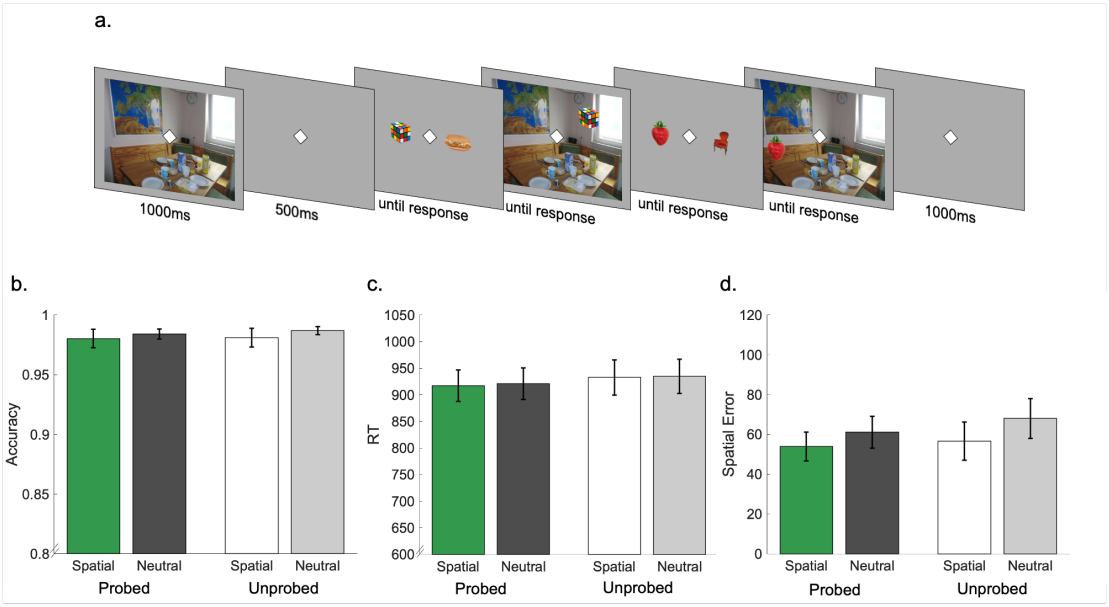


Figure 25. Panel a. illustrates the Retrieval task. Panels b-d show accuracy (in proportion), reaction time (in ms), and spatial error (in pixels) measures divided by whether the object in question had been probed and whether the context of the scene was paired with a spatially informative (green/white) or neutral retrocue (dark/light grey). There was only an effect of retrocue for spatial error (p = .003). Error bars represent standard error of the means.

Spatial errors when positioning the studied items within scenes showed that spatial memory was much better than chance (M = 427.66 pixels, SE = 1.83): Pairwise t-tests showed that the spatial errors for probed (spatial retrocue: t (39) = -19.53, p < .001, d = 3.09, neutral retrocue: t (39) = -18.44, p < .001, d = 2.92) and unprobed objects (spatial retrocue: t (39) = -18.54, p < .001, d = 2.93, neutral retrocue: t (39) = -15.81, p < .001, d = 2.50) were all significantly smaller than comparative chance values. Spatial

error was also assessed in a 2 x 2 repeated-measures ANOVA. Whilst there was no significant main effect of probe (F (1, 39) < .01, p = .946, η_p^2 < .01), there was a significant main effect of retrocue (F (1, 39) = 10.27, p = .003, η_p^2 = .21). Scenes associated with a spatially informative retrocue were associated with smaller spatial errors (M = 55.26, SE = 7.92) in contrast to scenes associated with neutral retrocues (M = 64.57, SE = 8.78), *Figure 25d*. There was no interaction (F (1, 39) = .27, p = .608, η_p^2 < .01).

Discussion

Study 2 successfully replicated the benefits of retrocueing for reaction time in a speeded recognition task. This replication suggests that the retrocueing effect for items in LTM is robust and worth building on. The observed benefit of retrocues at the increased tempo of the Retrocueing task in Study 2 also extend the findings of Study 1. Given that informative retrocues confer a behavioural RT benefit and show protection against incongruency costs despite the faster paste, it is more likely that proactive guidance is at play. It is especially noteworthy that informative retrocues appeared to shield against congruency effects not only for identification speed, but also for accuracy and spatial error. Indeed, the absence of incongruency costs following informative retrocues speaks to the idea of the proactive use of the retrocue, dealing with spatial information before an incongruent target location could slow or impair the response. However, it will still be worth exploring this with additional psychophysical as well as neural measurements. Spatial retrocues did not categorically improve localisation precision in comparison with performance after neutral retrocues (unless separated by congruency). Do retrocues thus only improve access rather than precision of the long-term memoranda? Would memory precision be affected if the retrocues provided

invalid rather than valid information, hence if the retrocue also encouraged participants to ignore rather than attend to the probed memorandum?

Study 2 did not replicate the effects of probing objects on subsequent memory. Participants were not faster for probed objects in comparison with unprobed ones. However, there was an effect of retrocues on localising the object in the scene in the Retrieval task: Spatial retrocues were associated with greater precision at later recall. It is possible that the differences in the Retrocueing task set-up in Study 2 changed any protracted effects of retrocueing and probing the remembered objects. Firstly, the increased cognitive demand of the speeded task conditions in Study 2 may have prevented participants from updating or refreshing their memory in a lasting way. This may explain why probed objects did not have a RT advantage over unprobed object in the Retrieval task. Secondly, the Retrocueing task in Study 2 implemented a localising task. Drawing attention to the location of the memorised items in the Retrocueing task may explain an effect on the later placement accuracy in the Retrieval task. Interestingly, the precision improvement was for items in scenes associated with a spatial retrocue, irrespective of whether the object had been probed or not. Spatial retrocues give information on which side of the scene the probed object is located. The localisation association of the unprobed object could have been boosted by inference that the unprobed object is on the opposite side of the scene to the probed object (indicated by the spatial retrocue).

Study 3: Comparing consequences of attending and ignoring items in LTM

Within the two previous studies, I observed a benefit in identification time following spatially informative retrocues in comparison with neutral retrocues that provided no predictive information. A similar pattern of effects occurred for slower (Study 1) and faster (Study 2) task conditions. These performance benefits were specific to reaction time and were tentatively attributed to selective attention operating within LTM to guide retrieval.

However, the comparison between trials with spatial and neutral trials may have introduced other unwanted variables. Firstly, the comparison of trials with informative and neutral retrocues may have introduced unintended differences in general state, potentially caused by arousal, anticipated task difficulty or need to process information. Arousal, for example, is known to influence performance in complex ways (cf Yerkes & Dodson, 1908). One might argue that the RT benefit of retrocues could at least partially be attributed to differences in state. Secondly, within Study 1 and 2, the retrocue could eliminate one (spatially informative retrocue) or none (neutral retrocue) of the potential matches of the forced-choice task: One could claim that in spatial retrocue trials, participants only had to match the one memory item to the probe, whilst in neutral retrocue trials both memory items were potential matches. The benefit in RTs for spatially informative retrocues might thus be a result of a difference in cognitive load.

Study 3 was designed to address these potential confounds. Rather than comparing trials with spatial and neutral retrocues, the current study compared performance when the retrocue accurately predicted the probed object (valid) with performance when the retrocue provided inaccurate information (invalid). Since the retrocue in Study 3 always provided spatially predictive information, any potential

difference in anticipatory state was eliminated. Secondly, as the retrocue was always informative, the cognitive load should be equivalent across conditions, at least within the preparatory stage.

Method

The experimental methods were approved by the University of Oxford Central University Research Ethics Committee (CUREC), and the study was conducted in accordance with the Declaration of Helsinki.

Participants

In total, 43 healthy participants (35 female, 8 male; 38 right-handed; age range: 18 to 34; M = 24.5; SD = 4.2) took part in our study due to overscheduling. Two participants had to be excluded, one due to an equipment failure and one due to an inability to perform one of the tasks. Hence the final analysis was performed on 41 participants (33 female, 8 male; 36 right-handed; age range: 18 to 34; M = 24.6; SD = 4.3). All participants had normal or corrected-to-normal vision and reported no difficulties with colour vision. They gave written informed consent before testing and were financially reimbursed for their time at a rate of £10 per hour.

Materials

The stimuli used were equivalent to the ones in Studies 1 and 2, but a greater number of scenes (80) and objects (440) were required in the main tasks given the nature of experimental conditions. Scenes and objects were still drawn randomly from the same sets of scenes and objects described above. All stimuli were presented using Cogent 2000 (RRID:SCR_015672), which ran via Matlab (2018, MathWorks, Inc.,

Natick, Massachusetts, United States). Study 3 was run on the same screens and under the same viewing conditions as Study 2.

Design

The tasks are illustrated in *Figure 26 – 28 a*. The same tasks were used as in Study 2 except for the following changes. A greater number of scenes were tested in the Learning and Retrieval tasks, which were otherwise identical. In the Retrocueing task (*Figure 8a*), all trials contained a spatial retrocue. However, the retrocue was only 75% valid, meaning that in the remaining 25% of the trials, the uncued object was probed instead (invalid retrocue). Prior to the task, participants were informed that the retrocue would indicate the object that was "most likely" to be probed. In order to ensure a sufficient number of trials in the invalid-retrocue condition, the overall number of scenes was increased to 80 (yielding 60 valid and 20 invalid retrocue trials). An additional modification was that the side of memory targets in the probe array was counterbalanced rather than randomised.

Procedure

The experimental protocol was equivalent to that in Study 2, with two exceptions. The larger number of scenes extended the study duration to approximately two and a half hours. The instructions for the Retrocueing task were updated to indicate the probabilistic nature of the retrocues.

Analysis

Data analysis for the Learning, Retrocueing, and Retrieval tasks followed a similar logic and approach as described for Study 2. The levels of the Retrocue condition, however, referred to valid vs. invalid trials.

Results

Learning task

Behavioural performance in the Learning task is summarised in *Figure 26*. The analysis of the Learning task replicated the findings from Studies 1 and 2: Participants were highly accurate in choosing the target objects and identification accuracy also increased systematically over blocks. In addition, their reaction times and spatial error decreased systematically over the three blocks, indicating good learning of object locations within scenes.

A 3 x 2 repeated-measures ANOVA on the accuracy data resulted in a main effect (F (2, 80) = 10.64, $p < .001$, $\eta_p^2 = .21$) and linear contrast (F (1, 40) = 17.91, $p < .001$, $\eta_p^2 = .31$) of block. Replicating the previous studies, accuracy increased systematically over the three blocks of learning (M, SE (in %): Block 1 = 98.37, .31; Block 2 = 98.77, .39; Block 3 = 99.18, .26), *Figure 26b*. There was no main effect of object (F (1, 40) = .61, $p = .439$, $\eta_p^2 = .02$), and no interaction (F (2, 80) = 1.24, $p = .295$, $\eta_p^2 = .03$).

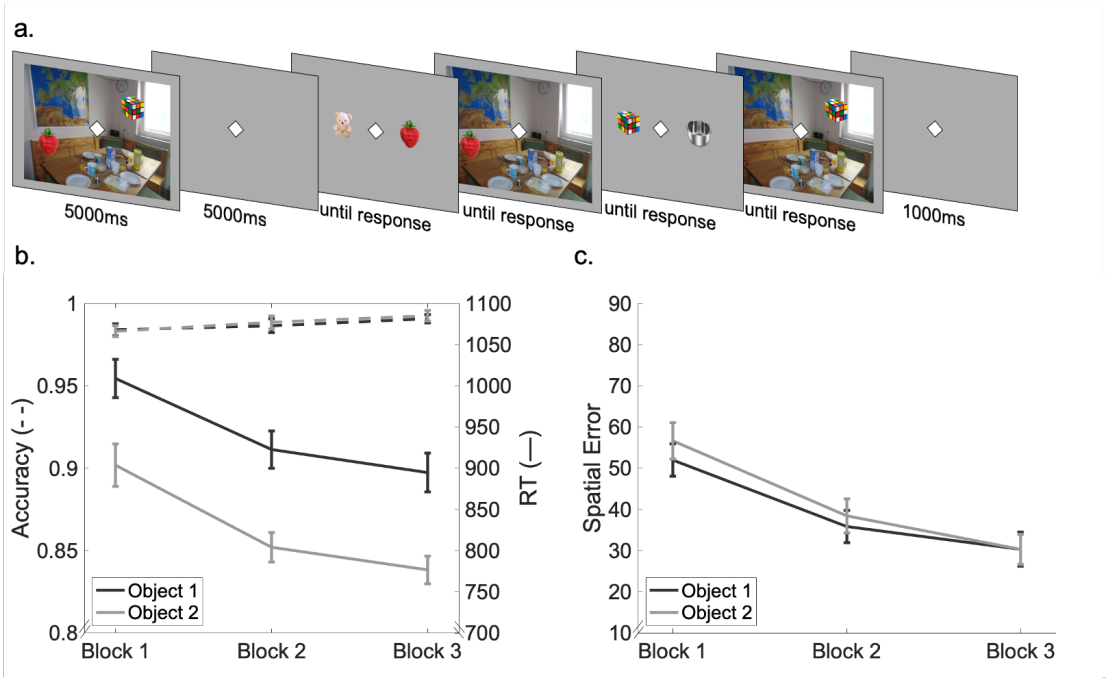


Figure 26. Panel a. illustrates the Learning task. Panel b. depicts the increasing identification accuracy (dotted lines, in proportion of accurately chosen objects) and decreasing reaction times (continuous lines, in ms) of participants over the three learning blocks. Dark grey illustrates these measures relative to the first probed object, and light grey to the second object. Error bars are standard error of the means. Panel c. shows the decreasing spatial error in pixels when positioning the two objects in the three learning blocks (chance level: $M = 432.80$, $SE = 1.59$).

A 3 x 2 repeated-measures ANOVA on the reaction-time data yielded significant main effects of block (F (1.52, 60.77) = 55.97, $p < .001$, $\eta_p^2 = .58$) and object (F (1, 40) = 48.38, $p < .001$, $\eta_p^2 = .55$). Reaction times decreased linearly over blocks (linear contrast: F (1, 40) = 68.07, $p < .001$, $\eta_p^2 = .63$; M, SE (in ms): Block 1 = 956.50, 22.37; Block 2 = 863.48, 18.42; Block 3 = 835.76, 17.31). Participants were generally faster to respond to the second probed object ($M = 828.20$, $SE = 18.86$) in comparison with the first object ($M = 942.29$, $SE = 20.96$), *Figure 26b*. There was no significant interaction between block and object (F (1.57, 62.74) = 1.02, $p = .352$, $\eta_p^2 = .03$).

Participants were accurate in their placement of the probed objects within the scenes. The average overall spatial error ($M = 40.61$ pixels, $SE = 3.78$) was much lower than chance ($M = 432.80$ pixels, $SE = 1.59$), t (40) = -28.97, $p < .001$, $d = 4.52$. A 3 x 2

repeated-measures ANOVA on spatial error data showed significant main effects of block (F (1.64, 65.58) = 118.85, p < .001, η_p^2 = .75) and object (F (1, 40) = 11.15, p = .002, η_p^2 = .22), *Figure 26c*. Spatial error values decreased systematically over blocks (linear contrast: F (1, 40) = 166.45, p < .001, η_p^2 = .81; *M*, *SE* (in pixels): Block 1 = 54.36, 4.13; Block 2 = 37.17, 3.99; Block 3 = 30.32, 3.88) and were lower for the first probed object (M = 39.40, SE = 3.68) in comparison to the second (M = 41.83, SE = 3.91). No interaction between block and object (F (2, 80) = 1.36, p = .262, η_p^2 = .03) was observed.

Retrocueing task

Replicating and extending Studies 1 and 2, analysis of the Retrocueing task showed significantly faster reaction times to recognise the memory target after a valid (M = 753.14 ms, SE = 30.88) than an invalid (M = 889.16 ms, SE = 42.19) spatial retrocue (t (40) = 6.06, p < .001, d = .95), *Figure 27c*.

In addition, identification accuracy and spatial error when positioning items in scene were also superior after valid compared to invalid spatial retrocues. A paired samples t-test showed that there was a significant difference between the retrocueing conditions for identification accuracy, t (40) = -2.06, p = .047, d = .32. The responses to valid retrocues were significantly more accurate (M = 90.07%, SE = 1.33) than for invalid retrocues (M = 84.45%, SE = 2.38), *Figure 27b*. Subsequent positioning of the memory target in the scene was much better than chance (M = 435.76 pixels, SE = 2.00) overall, and was significantly better after valid spatial retrocues (M = 70.69 pixels, SE = 9.32) compared to invalid retrocues (M = 132.54 pixels, SE = 19.53) (t (40) = 3.94, p < .001, d = .62), *Figure 27d*. Spatial errors were significantly lower than chance for both retrocueing conditions (valid: t (40) = -18.95, p < .001, d = 2.96, invalid: t (40) =

-10.37, $p < .001$, $d = 1.62$). No significant effects of congruency were observed for any of the behavioural measures.

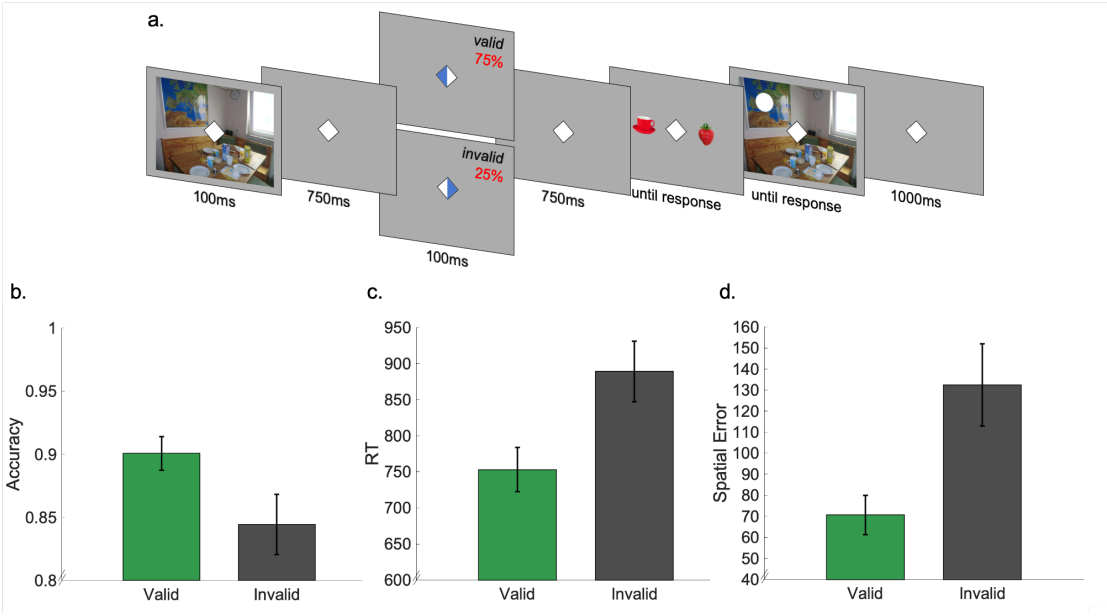


Figure 27. Panel a. illustrates the Retrocueing task. Panel b. shows the decreased identification accuracy for invalid retrocue (grey) trials relative to valid retrocue (green) trials. Error bars represent standard error of the means. Panels c. – d. demonstrate the decrease in reaction time (in ms) and spatial error (in pixels) in response to validly retrocued compared to invalidly retrocued objects.

Retrieval task

Performance in the Retrieval task was highly accurate across conditions with a mean of 98.31% (*SE* = .48%). A 2 x 2 repeated-measures ANOVA on accuracy data showed no main effect of probe (F (1, 40) < .01, p = .972, η_p^2 < .01) or retrocue condition (F (1, 40) = .41, p = .528, η_p^2 = .01). However, there was a significant interaction (F (1, 40) = 6.86, p = .012, η_p^2 = .15), *Figure 28b*. This interaction appeared to be driven by lower accuracy for objects probed after an invalid retrocue (M = 98.07%, *SE* = .60) and objects that were not probed after valid retrocue (M = 98.04%, *SE* = .55).

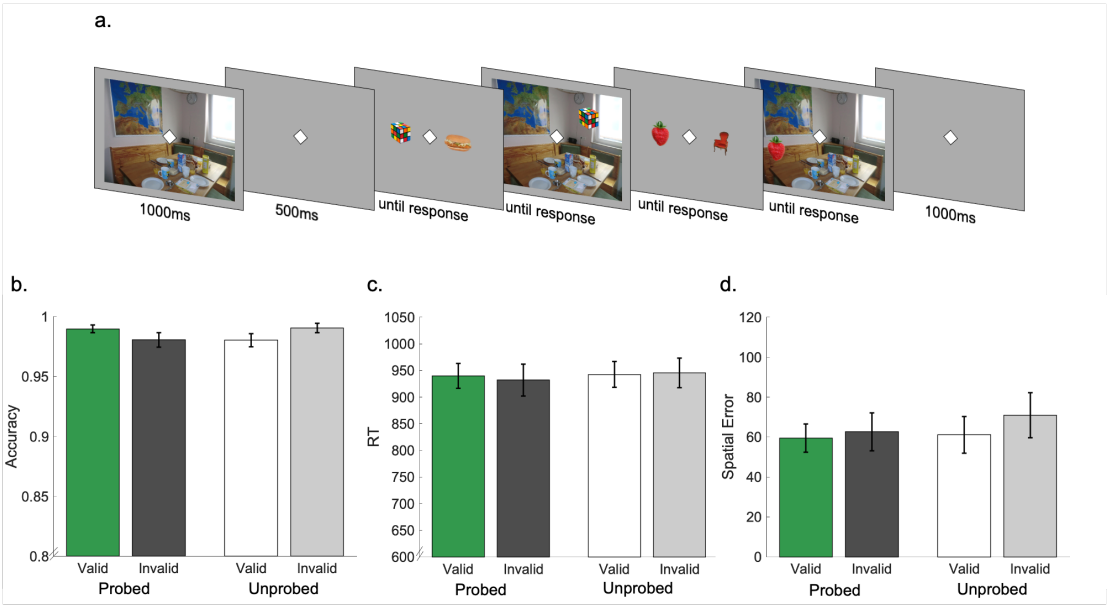


Figure 28. Panel a. illustrates the Retrieval task. Panel b. depicts the identification accuracy for choosing objects that had or had not been probed in the Retrocueing task and are within a scene context associated with a valid (green/white) or invalid retrocue (dark/light grey). Panels c-d. demonstrate that there was no difference in reaction time or spatial placement for objects, irrespective of probe and retrocueing condition. Error bars represent standard error of the means.

Reaction times for retrieving memory items were equivalent across conditions (Probe: F (1, 40) = .77, p = .385, η_p^2 = .02, Retrocue: F (1, 40) = .30, p = .586, η_p^2 < .01, Interaction: F (1, 40) = .50, p = .483, η_p^2 = .01), *Figure 28c*. Similarly, spatial errors were unaffected (Probe: F (1, 40) = .05, p = .822, η_p^2 < .01, Retrocue: F (1, 40) = .59, p = .447, η_p^2 = .01, Interaction: F (1, 40) = 2.52, p = .120, η_p^2 = .06), *Figure 28d*. Overall, spatial errors when positioning previously studied items within scenes were much lower than chance performance (M = 431.43, SE = 1.74), as revealed by pairwise t-tests comparing chance performance to probed (valid retrocue: t (40) = -18.50, p < .001, d = 2.89, invalid retrocue: t (40) = -17.08, p < .001, d = 2.67) and unprobed objects (valid retrocue: t (40) = -17.07, p < .001, d = 2.67, invalid retrocue: t (40) = -15.16, p < .001, d = 2.37).

Discussion

Study 3 is an important replication and extension of Study 1 and Study 2. Superior memory retrieval performance in the Retrocueing task shows the behavioural benefits are related to selective attention and cannot be attributed to other changes in general state or task load. In Study 3, the effects also extended to measures of accuracy and precision: Validly retrocued items were chosen more quickly and more accurately, as well as placed more precisely within the scene. These findings suggest that the effects of orienting attention within LTM are not limited to the access to a memorandum. The results of Study 3 suggest that selective retrieval can also influence the quality of a retrieved information. This warrants further study.

The direct comparison of valid versus invalid retrocue trials, without inclusion of a neutral-retrocue condition makes it impossible to separate attentional benefits and costs clearly. However, the apparently stronger and more widespread performance modulations in this task suggest that there may be costs of attending away from a memory item as well as benefits of attending toward an item. Further investigations are needed to characterise this effect in greater detail. It was not possible to pursue this within the constraints of this study, because of the escalation of experimental duration with such an increase in scenes and learning times. However, by using online training regimes, it should be possible to test participants on a larger number of scenes, and to add a neutral retro-cue condition additional to valid and invalid retrocues. A direct comparison between all three variants of retrocues would allow a direct judgement on the presumed benefit of valid retrocues and putative cost of invalid retrocues in comparison to neutral retrocues.

Any longer-term effects of retrocueing and probing on later retrieval were restricted to identification accuracy. Within the accuracy measure, the retrocue to an

object appears to give a memory boost to the object-scene association – irrespective of whether it was then probed or not. It is important to note that this relates to the memory association evoked by the retrocue: In trials with valid retrocues this corresponds to the probed object, and in trials with invalid cues it corresponds to the unprobed object (the farce for the later invalidly appearing probed object). The mental access rather than the physical act of the probe appears to lead to protracted effects on the memoranda.

Overall discussion

Over 3 studies, I was successful in developing an experimental approach that should be fruitful for investigating whether and how selective attention guides long-term memory retrieval. The core of the experimental approach is the Retrocueing task, which requires speeded recognition of memory items after different types of retrospectively informative cues. Due to the speeded nature of the task, the most sensitive dependent variable in the three versions of this task was reaction time. Participants were encouraged to respond as quickly as possible and the memory comparisons themselves were straightforward. The effects of speed of memory recognition were replicated across all three studies, which inspires confidence in both the effect itself and the ability of the task set-up to draw out the effect. The comparisons between performance after valid versus invalid spatial retrocues in Study 3 yielded further effects on identification accuracy and localisation errors. These results imply that selective attention might be able to optimise memory retrieval in multiple ways: Selective retrieval can both render the items in memory more accessible, and possibly, and at least temporarily, modify the representational quality of the memory itself. Further careful experimentation is needed to delineate these effects and tease apart under which conditions access or quality or both might be modulated by selective

attention. Another interesting finding was the absence of congruency effects as a potential consequence of this attentional selection that spatially informative retrocues conferred: The presumed prioritised state of the item selected by attention appeared to provide protection against costs of incongruent locations, both in terms of access and precision. Furthermore, the results hint that the information conveyed by spatial retrocues does not necessarily have a beneficial impact; it might also carry a cost if participants are encouraged to ignore the probed memory item. This suspicion needs to be confirmed in a further study including valid, invalid and neutral retrocues, which would allow a direct comparison of benefits (valid) and potential costs (invalid) of retrocues in comparison to a neutral condition.

The experimental approach developed across three studies was also successful in other respects: The Learning task elicited a highly reliable pattern of performance with high levels of learning and similar learning rates for different memory associations (in this case, objects in locations). Further, participants reported that the task was enjoyable and engaging: This presumably decreases levels of distraction and disengagement from the task ("just clicking"), which in my opinion can lead to unaccounted sources of confound. The simplicity and engaging nature of the learning paradigm also makes it an ideal candidate for adaptation for children or patient populations. The task could be modified easily to be used in online studies, simplifying the prospect of studying memory accuracy and the nature of long-term memoranda over different delays. This could include prolonged effects of any form of interference such as retrocues on memoranda. Any prolonged effects could then be assessed in the Retrieval task.

Indeed, within the current studies, the Retrieval task confirmed the high quality of learning from the Learning task. However, other results within the Retrieval task were more difficult to interpret due to potential interactions with the Retrocueing task:

Either the act of retrocueing or of probing the items in memory, or even both, could modulate later retrieval. Within the scope of these studies, it was difficult to tease these effects apart. Using a larger number of stimuli (scenes), and thus of trials, might give more power to these comparisons. Different ways to differentiate the effects of retrocue and probe could include assessments of the uncued item in the Retrocueing task, or the evaluation of only unprobed items (within the Retrocueing task) in the Retrieval task. The latter could evaluate whether the context (scene) was associated with a valid, neutral or invalid spatial retrocue. Moreover, it would be worth adding sensitivity to the Retrieval task to measure the effects of probing or retrocueing more accurately. Measures such as continuous delayed estimation tasks (compare Souza & Oberauer, 2016; Richter et al, 2016) or, more generally, task layouts that allow for mixture modelling may be desirable (Bays et al, 2009).

As discussed above, the differences in the results of the Retrieval task across the three studies might be due to the slight changes in the Retrocueing task in each study. The increase in difficulty due to speeded task conditions in Study 2 may explain the lack of a probing effect on later retrieval: The fast pace of the task may have prevented participants from processing the probed items fully, and to an extent that the object-scene association would be lastingly improved (cf. testing effect). Further in Study 2, the added localisation task – notably associated with longer display duration and increased RTs and hence possibly within the scope of a testing effect – may have drawn attention to the spatial location of the objects and hence influenced their precision at later recall. Interestingly, this effect was present within scenes associated with the spatially informative retrocue for both the probed and unprobed object. The nature of this effect can only be speculated about. Possibilities include inferential information about the unprobed object (opposite side of the screen from the probed object – as

specified by the spatial retrocue), and the reactivation of the second associated object in addition to the retrocued object before being selected against. The results of the Retrieval task in Study 3 might vary due to the very different nature of the Retrocueing paradigm: All retrocues were spatially informative, but the informative content could not always be trusted due to the presence of valid and invalid retrocues. In the following Retrieval task, identification accuracy was improved for items the retrocue was seemingly *predicting* rather than the items that actually *appeared*. For valid retrocues, this results in the same item, but for invalid retrocues this corresponds to the unprobed object. The Retrocueing task in Study 3 included the same fast pace as Study 2 with the increased response difficulty that invalid retrocues represent. Hence, the lack of lasting effects on the long-term memories on other measures might be attributed to a failure to process the probed associations deeply enough. This line of thinking encourages the idea that there might be different levels to which retrieval can be processed and that the nature of the retrieval matters on effects such as the testing effect discussed above (for example, recollection vs. recognition, see Greving & Richter, 2018). Such a possibility could also reconcile the diversity in findings on whether retrocues can improve only access to or also precision of WM memoranda: The task demands might enable participants to engage with the retrocued memorandum to a different extent. A comparison can also be drawn to the literature on suppression-induced forgetting (Anderson & Hanslmayr, 2014): Actively ignoring competing stimuli or even attempting to forget them might represent a much deeper level of interactive, selective retrieval than matching a retrocued item in WM or LTM. This might explain why I did not find the strong lasting effects of probing an object on later retrieval reported in the forgetting literature.

To understand the scope of how selective attention might operate in helping retrieval, various versions of the task and of retrocue types should be considered. Within the current set of studies, the retrocues gave predictive information about the spatial location of the probed memory items. Would other features such as colour constitute equally salient retrocues for LTM? Could even higher-order concepts (e.g., living versus non-living) or categories (e.g., type of object: toy versus kitchen utensil; or type of animal: bird versus fish) be implemented as retrocues to guide attention to part of a complex long-term memory? It would also be possible and interesting to manipulate the load of spatial/identity-object associations within scenes (or equivalent associations for non-spatial cues). This would enable the investigation of richer internal LTM search and the influence of informative retrocues on this process.

I have discussed considerations of access versus precision with regards to the probed memorandum, but can the precision of the retrocue have an impact on its own effect? Within these studies, the retrocue only cued one half of the scene. Would a precise (location or otherwise) retrocue have stronger effects? Are the attention effects on retrieval gradual or all-or-none? This will be worth investigating in further studies. A related question concerns what level of prioritisation a retrocue might bestow: Does retrospective attention to an aspect of a long-term memorandum (such as spatial location) activate other relevant concepts to that memorandum (such as colour)? It could be contemplated that retrocues to one feature of the memorandum bias attention to matches of other features of the same memorandum. Might these other features constitute guides or distractions under the right circumstances? A straightforward attempt to measure such attention effects beyond one feature might be as follows: Retrocueing a non-spatial feature of a memorandum paired with eyetracking could pick

up on gaze biases or even microsaccades towards the location of the retrocued item in question (van Ede et al, 2019).

In general, broader behavioural and neural measurements would be useful and desirable to understand the effects and underlying control mechanisms of selective attention in memory. Eyetracking might represent a simple way to assess proactive biases in a behavioural way, especially if the retrocues are spatial in nature and biases could be picked up via microsaccades or gaze biases (van Ede et al, 2019). Beyond that, electrophysiological measures will be helpful in confirming that the effects of selective attention indeed go beyond facilitation of matching during the recognition process and depend on proactive prioritisation of memory retrieval. Inspiration for potential modulations of attentional prioritisation can be drawn from retrocueing studies in WM and LTM studies studying memory-guided attention. Within the WM literature, retrocueing memory content has been associated with a decrease in memory load-dependent signals such as the contralateral delay activity in ERP (Kuo et al, 2012) and the desynchronization of the alpha-band oscillations attributed to shifts in attention (Wallis et al, 2015; Poch et al, 2014). These alpha lateralisations have also been observed in response to LTM guiding selective attention to goal-relevant regions in space (Summerfield et al, 2011; Stokes et al, 2012, Martín-Buro et al, 2020).

Functional neuroimaging techniques with more sensitive spatial resolution will be useful in outlining the brain regions executing this putative selective retrieval. One particular aspect of interest is to reveal the extent of overlap between selective attention to external stimulation, and within WM or LTM. Part of this enquiry would seek to evaluate the AToM theory described above (Ciaramelli et al, 2008). The mechanisms involved are likely to involve more than just the parietal lobe though, as retrocueing paradigms in the WM literature have noted additional activation of prefrontal regions

(Nobre et al, 2004; Nee & Jonides, 2009). Hence, to what extent does the dorsal frontoparietal network, known to control external goal-based spatial attention (Kastner & Ungerleider, 2000), also contribute to selective retrieval of LTM? It is important to note as well, that whilst the areas of interest might overlap extensively, their functionality might differ between selective attention and retrieval. Retrocueing studies in WM have found that the activation of the dorsal frontoparietal network common for selective attention to external and internal representation was only transient for retrocues, before being succeeded by activation in the cingulo-opercular network (Wallis et al, 2015), associated with action-related top-down control (Nelissen et al, 2013; Higo et al, 2011). Further, retrocueing studies in WM identified additional frontal areas particularly important for attention to internal rather than external space, such as the ventrolateral PFC (Nee & Jonides, 2009). Do these additional frontal areas also contribute to retrocueing in LTM? And are there specific contributions beyond the areas identified in WM retrocueing studies when LTMs are involved? Areas of interest for selective retrieval in LTM could be identified using the LTM literature. One such area of interest would be the hippocampus, known to have bidirectional connections to the parietal cortex (Cabeza et al, 2008): Parietal attention mechanisms could drive hippocampal retrieval processes, which in turn could trigger memory reactivation in the parietal cortex (for discussion see Tarder-Stoll et al, 2020).

Whilst the above-described neural measures will be useful to confirm the proactive prioritisation of long-term memoranda attributed to retrocues, our results yield tentative behavioural insight into the functional hypotheses of retrocues. The apparent efficacy of retrocues on stable LTMs provides evidence against a decay-based account of retrocue benefits. Pertinently, there is also no strong evidence for any refreshing or strengthening that the retrocues might bestow on the retrieved

memorandum, as I measured no clear protracted effects of retrocueing on the memory. Further, the observed effects of retrocueing on memory precision suggest that the benefit that retrocues bestow goes beyond mere facilitated access of a privileged state. However, it is important to note, that these precision effects only appeared with increased task difficulty and did not reliably affect the long-term stability of the memory.

The observation that retrocues affected precision under only certain circumstances also ties into the discussion of whether retrocueing effects will be task-dependent. The benefit of retrocues might be tailored specifically to what task demands are anticipated. Future studies will require using variations of the Retrocueing task beyond the addition of neural measures. Different task demands and stimulation parameters should help shed light on how attention modulates neural processing in the service of memory retrieval. It is quite likely that these attention effects are task-dependent, analogous to such task-dependent effects observed in WM retrocueing studies (Yu & Geng, 2019; Nobre & Stokes, 2019). However, how the nature of these potential modulations compares to attentional selection in WM remains to be seen. Our experimental approach could offer the opportunity to compare retrocueing in WM and in LTM within a common experimental framework. The task could be adapted to allow a fairly direct comparison between the types of memories, and hence to test correspondences in the neural and behavioural effects of attentional selection. One of the most interesting questions here also represents one of the greatest challenges: Are the effects on LTM actually buffered via WM and attentional effects on WM? As the encoding and retrieval of LTM could be argued to be processed via some instantiation of WM, it would be very difficult to tease the effects apart. Firstly, one might attempt to manipulate the amount of WM interaction the LTMs have, and thus get a graded

even if not absolute measure on whether such WM manipulations affect the long-term content of the memory. This would rely on the idea that such WM effects are additive. Secondly, the nature of attentional selection in WM could shed indirect light on the role of WM in LTM with appropriate long-term retrieval tasks. As described above, the effects of attention on WM can be task-dependent (e.g., Yu & Geng, 2019). If LTM and attentional effects on LTM are buffered via WM, it is conceivable that repeated instances of these distortions in WM might lead to biases in resulting LTMs. Feature-specific modulations could thus be tracked throughout WM and LTM tasks, providing a potential solution to tease WM and LTM effects apart across different conditions. Such task designs should be supported by electrophysiological and imaging measures that will help chart stages and types or processing that can be affected.

This might be an appropriate place to discuss that there are still open questions about this line of research. However, considering that this is only the first excursion into the possible efficacy of retrocues within long-term memory, it is unsurprising that not all issues could be addressed and resolved. The first of these considerations relates to the previously discussed topic on the relationship of WM and LTM. Specifically, there is an outstanding question on whether retrocues actually act on the retrieval of the LTM itself or on the re-instantiated representation of the LTM within working memory. This relates back to the idea of LTM as being "buffered" via WM (e.g., Fukuda & Woodman, 2017) or the idea of the activation of LTM supporting WM (e.g., Lewis-Peacock & Postle, 2008). With the present, purely behavioural measures, these possibilities might be very difficult to disentangle. Either way, it is worth stressing that the discussed studies present novel findings: Even if the retrocues would merely act on re-instantiations of LTM in WM, this is still the first time showing that retrocues can act on information that has been established for a long time, rather than just acquired

and modulated as is the case in usual WM retrocueing paradigms (e.g., see Souza et al, 2016). Further, at least arguably, the speeded conditions in the later studies make it more likely that the retrocue is acting during the retrieval process itself, both in terms of timing and in terms of pushing the perceptual limit. In other words, there would simply be very limited time for a participant to 1) retrieve the memory into a WM state and 2) use the retrocue to bias contents of the supposedly now working memory content. However, it would and will require further studies to disentangle the LTM and WM interpretations beyond doubt. One potential line of investigation would relate to the tracing the neurophysiological processes involved, which might be able to give indication on which level the retrocue acts on. Indeed, the pattern of neurophysiological activity in response to LTM retrieval that, for example, Fukuda and Woodman (2017) explore, could be used to trace the time point at which the retrocue shows efficacy. This could be used to dissociate whether the retrocue acts on the retrieval itself, or "only" once the memory is presumably held in WM. Alternatively, a different approach to disentangle these ideas might be to evaluate whether differences exist between the functionality and prolonged effects of retrocues in WM and LTM. Differences might allow to draw indirect conclusions about whether the cue acted on "just another WM" or indeed the retrieval of the LTM. There is already some tentative evidence supporting this line of argument: Whilst the pattern of protracted effects of the retrocues on the long-term memory within this study is rather opaque and does not follow a clearly interpretable pattern, retrocues in WM have been shown to affect the formation of LTM reliably, at least in young participants (Strunk et al, 2019).

A second, related point of discussion is regarding the retrieval itself: How would the proposed functionality of retrocues within long-term memory be any different from selective retrieval according to the selective retrieval literature (e.g., Wimber et al,

2015; Bäuml, 2019)? Before noting some crucial differences, I should expand on what the selective retrieval literature refers to. Selective retrieval as in the literature refers to the broader idea of retrieving a selected memory, as opposed to different possible ones. This is, for example, achieved by associating information (e.g., words, images, etc.) on two different levels: A with B, and A with C. A participant can then be asked to retrieve the "first association" selectively (e.g., Wimber et al, 2015). One might argue that such a task request could be considered a "temporal retrocue". However, while one might see broad similarities, in that the example utilises a retrospectively informative cue, there are some fundamental differences in the functionality of the cue, the memory material it acts upon and indeed the use of the term "selective retrieval". Considering the nature of the temporal cue, there appears to be a conceptual difference between cueing different time points within one continuous memory and pointing to different potential associations that happen to coincide with certain temporal information. Along this line of thought, I would argue that within the selective retrieval literature the selection process is not within one memory (as it is the case in these studies) but between different strands of memory that share a common denominator, or association: The memory component A is associated with several potential traces of memory, and what selective retrieval (in the literature sense) would do is activate one trace as opposed to the other, thereby weakening the other one. This weakening of the other traces is known as retrieval-induced forgetting (Anderson et al, 1994; Wimber et al, 2015). However, within the present studies and proposed mechanism of retrocues, the retrocue acts within one memory trace. This means, that the retrocue is designed to select and prioritise a certain aspect of this one memory trace rather than selecting amongst traces. There is thus a different understanding of the term "selective retrieval": In the literature sense, it was meant to select between memories, here it is meant to

select within a memory. Another way of phrasing this difference is with reference to cognitive processes: Retrocues are not thought to be a purely mnemonic process, but indeed reflect attentional processes within the memory. The present use of the term "selective retrieval" was coined in reference to the suggestion by Kuhl and Chun (2014) that the way we retrieve and work with memories largely represent acts of (selective) attention. The idea that retrocues rely on attentional processes is largely supported by the WM literature with one hypothesis stating that the attentionally selected object within WM gets afforded a prioritised state and transformed into a "*prospective, action-oriented representation*" (Myers et al, 2017). Further, WM studies using dual retrocues support the idea that retrocues manipulate the attentional selection of items within memory, allowing the currently selected mnemonic item to be decodable (LaRocque et al, 2013). Given that the outline of the current paradigm consciously mirrored these WM retrocueing studies, it is not far-fetched to assume that such attentional processes would be at play, even though they were not measured directly via neural or electrophysiological measures. The current results do bear the hallmark of attentional processes, showing faster reaction times and more accurate identification in response to valid, informative cues as opposed to invalid ones. Therefore, the etymology of the term "selected retrieval" appears to differ somewhat between the contexts. As opposed to the selective retrieval literature, the memory cue refers to the entire, not part of the memory, meaning presumably that the entire memory context (in this case including both object identities and their locations) are getting retrieved. During this retrieval process, attentional direction via retrocues putatively highlights aspects of the memory that then get preferentially retrieved. This hypothesised process characterises the novelty of the results, contrasting it with the broader understanding of selective retrieval in the literature. Indeed, the observed differential effects of the retrieval itself support

the hypothesis that they rely on different processes: Whilst retrieval-induced forgetting is a well-known consequence of selecting some memory associations over others (Anderson et al, 1994; Wimber et al, 2015), within the current set of studies I have observed no indication of forgetting the non-retrieved object in the subsequent memory task. Of course, one would require more studies delineating this differentiation, but I believe the evidence presented within these studies, and the known effects of retrocues on WM and, separately, retrieval induced forgetting all provide convincing evidence for this distinction.

A third topic considers potential differences of selective retrieval by cueing attention prior to memory retrieval (precues) or after initiated retrieval (retrocues), if the memory in question is already established. Within the working memory literature, there is a fundamental difference in the cues by value that the mnemonic content is known in one case (retrocues) and yet unknown in the other case (precues). WM literature supports benefits to the memory by either type of cue, even though precision benefits were greater in response to precues (e.g., Wallis et al, 2015). Wallis and colleagues argue that the cues might act on different gating processes postulated by Hazy and colleagues (2007), with precues facilitating input and retrocues facilitating output gating. However, what would be expected if the cues act on an established long-term memory? The "pre" of the precue in this case would refer to a cue prior to retrieval of the memory. Conceptually, there still appears to be a difference, in that a precue in this instance would imply that one is unaware which memory or aspects of the memory should or will be selected and retrieved, whereas for a retrocue the retrieval process is presumably already initiated. Considering the example of a traumatic flashback, the difference between the two becomes striking: A precue, suggesting which part of a traumatic memory to focus on and select, might be a useful tool when consciously

retrieving trauma memories for example in a counselling setting. If, however, a traumatic memory comes to mind involuntarily, as is the case for flashbacks, any influence on the retrieval process would have to be retroactive. The hypothesised ability to influence retrospectively which aspects of the memory are going to be selectively retrieved once the retrieval is already initiated might open an entirely different way of dealing with trauma memories. If retrocues are indeed effective in the directing selective attention during the retrieval process of long-term memories, as is suggested by this set of studies, it might allow modifying which aspects of the (involuntary) trauma memory are selectively retrieved. Of course, it would be interesting to compare the effects of precues and retrocues on long-term memory by contrasting them within one task design. However, it should be noted that both types of cue might merely lead to behavioural enhancement of the attentionally cued mnemonic content, as has been observed in respective WM studies (e.g., Wallis et al, 2015). It might require careful experimental manipulation to tease potential differences in behavioural effects apart, for example by comparing effects of the cueing on longer term memory or on congruency. Congruency might be a promising route given the observed results: Without going into the literature on congruency or Simon effects, one might expect that advance information about the memory to be retrieved to leave the congruency effect intact (e.g., Wascher & Wolber, 2004), whilst the informative retrocue in this set of studies practically eliminated any congruency effect. Either way, the neural processes supporting these pre- and retrocueing processes can probably be contrasted more easily than pure behavioural measures, for example by tracking the time course of the preparation and response to the memory with EEG.

As revealed by this discussion section, the three presented studies provide strong and arguably the first evidence for the idea that selective attention could also serve a function in the retrieval of stable and complex long-term memories. The design using retrocues represents a promising experimental foundation for further exploration and adaptation to various behavioural and neural measurements. Our observations take the first early steps into the fascinating investigation of how prioritisation and selection processes operate to guide effective and adaptive retrieval from long-term memories.

Chapter V. General discussion

The work in this doctoral thesis explores mutual influences between memory and attention at the three main phases of memory establishment, namely: encoding, consolidation, and retrieval. At which of these stages is the memory malleable enough for attentional modulation? Could such attentional impact influence the emerging memory and its subsequent utility in guiding attention? Within the three experimental chapters, I have attempted to answer these questions in relation to encoding (Chapter II), consolidation (Chapter III), and retrieval (Chapter IV). Each chapter discusses the nature and results of the studies within the context of their specific scientific background and motivations. However, it is worth considering the implications of each body of work within the broader scope of memory and attention, finding a niche for each within the study of memory-attention interactions. Seen in the light of the relevant literature, some of my doctoral work highlights particular elements of attention and memory, thereby sparking novel research ideas and inspiring some further, sometimes admittedly speculative, research hypotheses. Moreover, I wanted to give some space in this discussion to issues beyond the scientific results, the "meta-findings" as you will. These might consist of methodological problems I encountered and the development of potential solutions, lessons learnt about planning and conducting research, or assessing the utility of collected data. For the sake of readability, the discussion points are structured along the three strands of research and memory stages discussed.

Distraction during memory encoding

The study in Chapter II examined the consequences of manipulating distraction during the encoding of spatial contextual associations on subsequent memory

performance and its use as an attentional guide. Specifically, I was motivated by the idea of proactive suppression: can we learn to suppress the presence of a distractor over time, and then use this learnt distractor information to inhibit its associated location? The putatively encoded "distractor suppression" was traced behaviourally and using electrophysiology along two lines of investigation: Firstly, I examined the effects of distractor suppression on memory formation and on its utility in guiding attention, contrasted with memories for which encoding was undisturbed by the distractor. Secondly, I attempted to measure whether an inhibitory tag had formed at the distractor-associated location.

Interestingly, despite the confirmed distracting effect of the salient distractor stimulus, participants were nevertheless able to complete the search task and form memories of the target-location associations in the face of significant distraction. The encoded target memories persisted, and participants showed a good meta-awareness of the quality of their memory, irrespective of distractor association. Overall, the encoded memories did not show a predominant impact of the distractor on the quality of the memory nor its utility as a guide for spatial attention. The divergent observations between significant distraction effects during the encoding process and the lack of strong distractor-based consequences for memory retrieval or attentional guidance were only made possible by separating the learning and testing phases in the study. Such separation was not present in other studies (e.g., Noonan et al, 2016), reflecting positively on the design of the present study. In general, the distraction study reveals a notable level of robustness in the encoding processes, withstanding significant distraction for the benefit of retaining precise information for future use.

Whilst the overall, global memory quality was left unperturbed by the distractor, some modulation occurred in the direct vicinity of the distractor: behaviourally, the

reported target location was biased towards the distractor location. Could this be some evidence for the development of an inhibitory tag the study was designed to test? The spatial distortion indicated at least some residual trace and influence of the distractor association. Several alternative mechanisms could have resulted in a systematic bias of memory reports toward the distractor location. Along with the hypothesised inhibitory tag, the distractor location could have been inhibited, leading to a warping of the encoded space. This warping effect in turn could have resulted in a spatial bias towards the epicentre of this distortion, the distractor location itself. Further, an attentional pull towards the distractor location could be the result of residual attentional elevation associated with the distractor. Alternatively, learning the location of a salient and task-relevant item (a distractor that had to be suppressed) could have resulted in enhancement of the distractor location that subsequently attracted attention to that location. If the distractor was registered as a task-relevant item, it could contaminate the report of the target location. Another explanation is the use of the distractor as an anchor for relational learning: if targets were remembered in relation to the very salient, second contextual landmark that is the distractor, it might result in the representation of both objects drawing closer in conceptual and spatial distance.

It is also difficult to rule out that participants may have tended to report locations closer to the centre of the screen ("centre bias"), leading to distorted reports that would approximate a pull toward the distractor location. The latter explanation would present a problem for most studies reliant on explicit spatial recall in a two-dimensional environment and might only be somewhat addressed experimentally by using large or three-dimensional spaces in experimental designs, such as virtual reality (VR), in which distractor locations are equally likely to be further or closer to central locations. Within the context of this specific study, an easier approach to disentangle some of the potential

explanations for the putative distractor pull would be to add an additional an explicit recall task for the distractor location. It would be necessary to add this task at the very end of the experiment rather than intersperse it with the present memory task. A separate task would prevent an elevation of the distractor in mnemonic relevance, especially cumulatively towards the end of the memory task. In general, the observations of more local distractor effects are reminiscent of distinctions in the learning and use of global and local contexts in contextual cueing paradigms (see Brockmole et al, 2006). It would be worth pursuing the possibility of global and local effects of distractor association on memory in their own right.

Although the behavioural data showed little impact of the distractor association on the ability to orient attention to target locations based on memory, the electrophysiological data were more suggestive of an influence. The distraction-related modulation of the ERP signal in response to the scene cue suggests that the distractor association was reinstated and had to be dealt with at least to a certain degree. For this to occur, mnemonic information about the distractor had to be retained differentially from the remaining memory context: even though participants were unable to suppress the distractor during encoding and any potential distractor effects remained undetectable in behavioural measures, distraction information must have been incorporated and processed proactively preceding target appearance. Indeed, the pattern of electrophysiological signals in response to the target appearance supports this conclusion: The presence of distraction during encoding amplified the difference in neural activity related to processing items appearing at expected (valid) vs. unexpected (invalid) locations. Is it possible that memories forged in the face of distraction are better defined or less flexible in their use? It is conceivable that the connections between target location and scenes created in circumstances of higher attentional-selectivity

demands due to a distractor might also be more selective in their use, causing a greater differentiation between conditions: Expected target locations could be enhanced more and unexpected locations could be further deprioritised. There is evidence supporting the idea that memory is strengthened when encoded with concurrent, successful suppression in an inhibition-of-return scenario (Markant & Amso, 2013; Markant et al, 2015). The data within this study are not conclusive on these theories, especially considering the absence of successful distractor suppression during encoding and of a behavioural distraction effect during memory-guided attention. It would be interesting to pursue these ideas further in a study tailored to achieve significant and well documented distractor suppression during learning.

The considerations about an inhibitory tag at the distractor location follow a similar direction: whilst I did not find a significant behavioural cost in responses to items appearing at the previous distractor location, subtle but significant differences arose in electrophysiological measures. Did this putative inhibitory tag start to form, yet was not powerful enough to be expressed behaviourally? Would this have changed if participants were successful in learning to suppress the distractor during encoding? Only further studies could tell.

Interestingly, although the electrophysiological data point to a residual consequence of the learned distractor location, they do not clarify the nature of its effect. It remained difficult to determine whether the learned distractor location exerted an inhibitory effect or whether the learned distractor led to an excitatory effect related to the learning of a second task-relevant association. The neural trace appears to "match" responses to the target location initially, before diverting to "mirror" responses to other invalid locations. Does this pattern of modulation provide an indication of why behavioural measures appeared to show responses to items at distractor locations as

intermediary between those to items appearing at valid and invalid ones? Or might there be an underlying difference in how participants treated the distractor and thus its associated location? As the memory for the distractor location was never probed in the participants, it is entirely possible that at least some participants treated it as a second target. For example, some participants may have assumed that the location of that salient item was likely to be probed despite the experimental instructions. Others may have followed the instructions more rigorously and actively tried to ignore or suppress the salient distractor during learning. Naturally, this variance might be washed out behaviourally when averaging across individuals and only be retained in the time-sensitive ERP traces if the processing steps occurred within different time windows. It might be prudent to address this possibility experimentally by explicitly referring to the distractor as such, or by creating a between-subject condition in which the approach to the distractor is varied. In any case, the results show that despite the absence of a behavioural effect of the distractor on memory-guided attention, the presence and location of the distractor during encoding did leave sufficient mnemonic traces to influence subsequent processing of events appearing at the distractor-associated location within a scene.

A recurring theme when discussing the distraction study is the difficulty in measuring suppression-related effects, given that suppression did not measurably occur during the encoding process. As such, the study is a lesson in the importance of thoroughly designing each part of a study, especially given the interrelated nature of studying memory and attention interactions. Firstly, the nature of the distractor was an important factor. A pilot study had suggested that a stationary distractor might not be sufficiently attention grabbing, especially given the visual complexity of the rest of the scene. However, the chosen distractor might have been too successful, preventing

participants from forming an inhibitory tag. It would have been useful to run a pre-study manipulating the strength of the distractor stimulus and evaluating participants' ability to suppress different levels of distraction, as well as to measure how much exposure is necessary to achieve a reliable effect. Further, the flickering appearance of the distractor, which was motivated by findings of oscillatory entrainment to remembered locations (e.g., Wimber et al, 2012), ended up interfering with electrophysiological assessments of preparatory attention, as the frequency of the flickering overlapped with the alpha band. If a distractor should flicker, it is important to achieve a distinct frequency that does not overlap with an oscillatory frequency of interest.

Secondly, in line with evaluating the need for more exposure to the distractor to achieve suppression, the study would have benefited from a more successful encoding regimen. The high exclusion rates due to memory failure suggest that participants may have needed more learning blocks to encode the vast breadth of information lastingly. All other behavioural and neural effects may have been diminished as a result of an insufficiently strong memory base: with higher memory precision, more subtle effects of distraction might have been discernible. Moreover, its strength and precision might well be a factor when using the memory as attentional guide. The variance based on differing levels of memory guidance could have veiled potential distractor effects. Within the limits of this study it was not feasible to include more learning blocks and thus better encoding, but the observation emphasises the importance of developing successful encoding tasks to ensure successful memory formation at less cost to experimenter and participants. The learning paradigm developed in the study on selective attention during retrieval (Chapter IV) might represent such an option: An interactive design allowed successful memory formation of complex, varied

information over fewer trials. Indeed, learning tasks can and should be informed by what research has shown to improve learning and the retention of information, as encoding is a necessary precursor to studying further effects on memory down the line.

Returning to the topics the distraction study sought to illuminate, I can conclude that the encoding processes can withstand a substantial amount of attentional distraction in support of forming lasting memories, whilst preserving their utility as attentional guides. However, this does not mean that the encoding process was impervious to distraction-related attentional modulation, as evidenced by distortions within the spatial relations of the memory and differential electrophysiological processing in response to the memory and its components. In parallel, I found no clear evidence for the proactive use of mnemonic information about distractors in the guidance of attention. Given the discussed limitations of the setup, this does not imply that it is impossible to form such an inhibitory tag. As it happens, the differential signature of the electrophysiology and in particular, the time-course of this modulation, suggests that inhibitory processes may have occurred but, being overcome rapidly, were not pronounced enough to affect behaviour. However, any strong conclusion will necessitate additional studies, preferably with reliable evidence to distraction and distraction suppression during encoding.

Targeted memory reactivation during memory consolidation

The second project, the "nap study" (Chapter III), considered the consolidation period: Can we influence which memories are consolidated, and does this manipulation have knock-on effects on the emerging memory and its utility as an attentional guide? In this study, I attempted to bias the consolidation of some memories over others and, in this way, approximate a selective attention manipulation during consolidation which

might render those selectively encoded memories more potent. The "attentional" manipulation of the consolidation phase was attempted by targeted memory reactivation (e.g., Oudiette & Paller, 2013) during a napping period: During deep sleep, I replayed tunes to the participants that were associated with half of the encoded memories during a previous learning phase. Specifically, I was interested in whether the memories driving memory-guided attention were subject to sleep-based consolidation. Further, the manipulation probed whether an auditory context, which was conceptually unrelated to the scenes, can trigger replay during sleep and thus prioritise memory consolidation. Lastly, I wondered whether sleep-related effects on memory, predominantly observed to be benefits, would also impact on their utility as attentional guides.

Fortunately, the basic foundations of the study set-up proved to be suitable: Participants successfully encoded and retained the target location-scene associations of the learning task irrespective of the additional pairing with music. Even though the music level added complexity to the mnemonic information, participants retained accurate representations of the memory both before and after the nap opportunity, as well as having an accurate impression about their own performance ("meta-awareness"). In the present study, I did not explicitly test whether a memory association between scene type and tune had formed, but this could be achieved by adding an explicit recall task or a recognition task: A tune could be played in parallel to a scene showing, with participants having to indicate whether the scene was associated with this tune or not. In addition to showing the usual effects of the learning and memory tasks, the study replicated the standard behavioural memory-guided attention effects (e.g., Summerfield et al, 2006). Further, the design of the nap opportunity was successful, allowing the majority of participants to fall asleep deeply and for some time.

There was also a high level of agreement across sleep scores for the polysomnography data, especially for the sleep stages of relevance. These favourable outcomes allowed investigating the main objectives of the study outlined above.

The first line of investigation was concerned with the effects of sleep on memory and its utility for guiding attention, both of which appeared to be very limited at best. Contrary to many findings reported in the literature, I did not observe a big qualitative difference in memory as a consequence of taking a nap. As the measures of interest involve averages of the memory performance, it is entirely possible that memories for individual scenes might have benefited from sleep or suffered from the delay. However, there was no evidence for a systematic effect within and across participants. Participants were faster to respond in the memory session after their nap, yet there are multiple explanations that could account for this finding, including the positioning within a testing session (at the end of a long session vs. within a very short session), level of familiarity with the general task procedures, or the time of day. In the context of the literature with well-established benefits of sleep on memory consolidation, the absence of sleep-related benefits in this study was surprising. It is possible that any benefits were masked by detriments of having a delay between the nap and the second testing session, or by greater confounding effects of mnemonic inaccuracies caused by insufficiently strong encoding. The observation in the literature on the importance of perceived future relevance for consolidation (Wilhelm et al, 2011b) might also explain the lack of sleep-based effects: Participants were unaware of the subsequent memory task after the nap opportunity, which might have discouraged them from "paying attention" to the memories and thus, putatively, their consolidation.

A further surprise, given the literature, was the observation that longer sleep durations were associated with higher spatial errors during spatial memory retrieval.

Sleep duration, however, was not associated with performance in the memory-guided attention task. And indeed, the detrimental association between prolonged sleep duration and memory quality did not persist when the interpersonal variance was removed, achieved by measuring and normalising forgetting across memory sessions. Taking these three aspects into account, along with the well-known saying that "correlation does not imply causation", it might be worth considering the findings not in the light of sleep duration but rather in light of who slept: Participants who slept long, may have just been in need of more sleep and thus performed worse on the cognitive tasks. Including a control group without nap opportunity, whilst not remedying the concern, might alleviate the implications of this consideration. Overall, I did not find support for the idea that the memories driving memory-guided attention are impacted by sleep-based consolidation. However, given the lack of effects on the memory itself, this set of data is unsuited for drawing any meaningful conclusions on this issue.

The lack of selective targeted replay effects on memory prompted the second line of investigation: the adequacy of an auditory context as a cue for biasing memory consolidation. As we did not observe a categorical influence of stimulation on memory quality, or indeed memory-guided attention, the results could be interpreted as evidence for the unsuitability of unrelated auditory cues for targeted memory reactivation. However, the measures I used for the categorical contrast are fairly crude and, in addition to multiple shortcomings of the study, might not accurately reflect all processes.

Firstly, an unintended but major problem for this study was the low number of participants. Traditionally, nap studies, and in particular studies using targeted memory reactivations, must adhere to strict exclusion criteria, often excluding about half of the tested participants (e.g., see Westerberg et al, 2015; Vargas et al, 2019; Wang et al,

2019). The unusual levels of exclusion in these studies are due to having to ensure that participants fall asleep deeply and long enough to ensure a significant amount of stimulation. Additionally, memory-guided attention tasks can end up having high levels of exclusion (see Chapter II), due to the need to establish a memory before one can argue that it guides attention. If participants do not establish enough of these memories, the trial numbers are too low to support credible results. Given uncontrollable external circumstances (one notable circumstance being the global pandemic), it was impossible to "top-up" the sample to the required size, once the extent of the required exclusions became apparent. It was thus necessary to strike a difficult balance between applying reasonable exclusions that would ensure that participants had formed memories, slept, and received stimulation, whilst at the same time retaining a defensible sample size. Future studies could remedy this by developing simple measures to assess the viability of a data set for inclusion, ensuring that the required or at least intended sample size can be matched during the data collection phase.

The second consideration concerns the auditory tunes themselves. The complexity of the tunes, whilst potentially riskier from an experimental point of view, should not have precluded targeted memory reactivation: Antony and colleagues (2012) were successful in cueing targeted memory consolidation for skill learning by using whole melodies. Instead, the difficulty may have lain in in the strength and relevance of the association between scene and tune, which may simply have been too weak. Informal discussions with researchers in the field suggest that conceptually unrelated associations may need to be over-trained to be effective as cues. Relatedly, it was not possible to play the entire tune during sleep, which may have resulted in replayed tunes not having been well registered as re-instantiations of the previously experienced auditory context. Even if the entire tunes could be used, one would still require fades

in and out to prevent arousal, just like in the present study, which in turn could change the recognition value of the tune. This might be especially important given that the fast reaction times in later learning blocks imply that only the beginning of the tune was heard. As discussed in Chapter III, it is yet unclear which part of an auditory cue is essential for its efficacy in targeting memory reactivation. Is it the beginning? The most recognisable part – which might be subjective? Either way, a strong, almost exaggerated association between scene and tune might have made them, and potentially parts of them, more effective as retrieval cues.

The third problem also concerns the auditory cues, specifically their match to a category of scenes, rather than to individual scene contexts. The mapping between tune and scene was therefore one to many rather than one to one, as is most often used. During the design phase of the study, arguments for and against having a separate tune for each scene were contemplated and weighed. A strong argument for a single pairing is obviously the clarity of creating unique episodes in which an item is associated within a spatial and auditory context. Such a design would also match other studies in the field. However, the convincing argument against the one-to-one pairing was the sheer memory load on the participants. Given that it is already difficult to ensure that participants form stable representations of a location in space for the number of scenes required to satisfy all conditions, how could they also remember an individual tune for them? How was it possible to ensure that participants with an untrained musical ear could distinguish between all these tunes? In other targeted memory reactivation studies with auditory cues, the cues are inherently linked to the object (e.g., a meow for a cat, Oudiette et al, 2013), which does not create a separate level of information to be linked and retained. I thus settled on having four different categories, matched with four different tunes, in the hope that the auditory context would reactivate scenes of the

associated category. However, the memory load may have simply been too great to cause categorical, reliable effects of replay. Further, as mentioned before, the categorical reactivation of multiple conflicting memories could have interfered with beneficial memory consolidation.

A fourth point of reflection is not as much a problem for the current study, but rather an inspiration for further study designs: Considering the argument that targeted memory reactivation was reliant on the presence of REM-sleep (Tamminen et al, 2017), it is possible that the full effect of the stimulation might only be measurable after a longer sleep opportunity, such as a night's sleep. Even without the possibility of conducting overnight sleep studies, this issue could be addressed by adding a behavioural testing session the next day, mimicking the post-nap session. Should targeted memory reactivation during slow wave sleep indeed predominantly serve to give memories an "attentional tag" for later consolidation, then overnight sleep might allow the biased consolidation to take its full effect. This consideration is especially interesting given the observation of a potential locus or functional underpinning for such attentional tagging in animal studies of the thalamus (see Klinzing et al, 2019; Chapter I, Attentional Influences on memory consolidation).

Lastly, a lack of replay effect might be due to the limited amount of stimulation that occurred. This is related to the point raised above about the number of participants: There were participants who had received multiple rounds of stimulation, even if not in comparable numbers to other studies (e.g., see Schechtman et al, 2020), yet they represented a small subset of participants. Indeed, analyses on the duration of received stimulation suggests that the amount, or rather duration, of stimulation given might have been of crucial significance. Overall, stimulation duration did not show pervasive correlations with memory performance, or even forgetting. However, when only

considering participants who had received higher levels of stimulation, a beneficial effect emerged: Higher stimulation duration was associated with less forgetting, indicating a protective effect of prolonged stimulation during sleep. It is thus entirely plausible that the lack of memory-related effect of the targeted memory reactivation in the present study was due to insufficient stimulation duration in the majority of participants. The logical continuation of this explanation would be to consider the effects of stimulation as cumulative, observing generally stronger patterns of results for the subset of participants with more stimulation. Notably, the process might still require a certain threshold to be reached before the effects become pronounced. Further, the proposed cumulative nature might not be linear, increasing consolidation progressively by stimulation, but might occur stepwise: With each stimulation, more memories might have had the chance to be reactivated and consolidated, thus on average continuously increasing precision over all memories. It would be interesting to explore this possibility in a structured, scene-by-scene analysis given suitable data.

The most surprising result of the study by far was the strong negative relationship between the duration of sleep-based stimulation and the utility of memories in guiding attention. This analysis focused on participants with a minimum amount of stimulation, so it is important to be mindful of the results representing a smaller sample size. Nevertheless, the pattern was striking: The more stimulation participants received, the less they benefited from informative mnemonic information in contrast to misleading information. Further, the more stimulation they received, the more the utility of "replayed" memories decreased all together in comparison with scenes that were not biased for consolidation. These findings were antithetical to my predictions and I have discussed possible explanations within Chapter III. However, if I can indulge in speculation, this observation might have uncovered something interesting about the

functional nature of memory-guided attention. Again, I have to stress that this is speculation, and the current sample size by no means allows decisive conclusions. Keeping these limitations in mind, bear with me: It has been shown that sleep-based consolidation changes the quality of memory (see Born & Wilhelm, 2012): Memories become more abstract, like schemata (Lewis & Durrant, 2011), often moving from more specific to gist-like representations (e.g., see Moscovitch et al, 2016; Klinzing et al, 2019). These gist-like representations are less hippocampus dependent (Winocur et al, 2010). Biasing some memories for consolidation over others might render the selected memories less specific, more abstract or gist-like. This process might decrease the hippocampal dependency of the memories, or at least the hippocampal involvement when they are retrieved or used. Given the emphasis in the literature on the spatial nature of the hippocampus (O'Keefe & Nadel, 1978; Burgess et al, 2002; Buffalo, 2015) and the importance of the hippocampus for attentional guidance (Summerfield et al, 2006; Goldfarb et al, 2016; Günseli & Aly, 2020), perhaps the memories targeted for consolidation do not benefit as much from the hippocampal-mediated attentional orienting because they rely less on hippocampal input than their counterparts. Again, this is speculation. Still, the thought experiment might inspire some approaches to investigate the functional nature of memory-guided attention. The hypothesis would predict that decreasing hippocampal dependence would be accompanied by decreasing flexibility of the memory and decreased utility in guiding attention. Evaluating this claim would require using a neuroimaging technique sensitive to hippocampal involvement, such as fMRI. Further, it offers a methodological approach of tracing such a development over time: the utility of the memory in guiding attention could be contrasted after multiple, varying delays, including sleep-based consolidation with or without employing the technique of targeted memory reactivation.

Speculation aside, the current study did provide some evidence for the effects of targeted stimulation during sleep on memory-guided attention. Whilst there was no clear, homogenous effect of stimulation on memory and attentional guidance, a more nuanced picture emerged when the stimulation duration was taken into account. Longer stimulation durations showed a limited protective effect on forgetting, and a detrimental effect on the utility of the memory in guiding attention, thereby indicating the presence of some "attentional" modulation of the memories during consolidation. This effect remains to be replicated with an appropriate sample size, and the underlying functionality driving the effect can so far only be speculated about.

Beyond the scope of the results of this study, there are broader considerations about the nature and use of sleep studies. Whilst larger sample sizes might afford the ability of more detailed or specific analyses, they are bound by the set-up of the study. For example, the analysis of overall oscillatory power in different frequency bands proves to be difficult if complex auditory stimuli are played during the sleep period. This can present a problem for all targeted memory reactivation studies but is especially conspicuous with complex stimuli: the auditory processing of the tunes will likely interfere with straightforward measures of oscillatory power. This holds especially true as the tunes, for counterbalancing reasons, will likely differ between participants, and the frequency with which the tunes are played will differ due to the irregularities in sleep behaviour. Potentially triggered differences in the oscillatory pattern might end up clouding or distorting any differences due to replay. The possibility of running such analyses is thus dependent on a carefully balanced experimental protocol during sleep, such as introducing set windows with and without stimulation, that might serve as basis for comparisons. Further, the use of theory-driven analysis strategies such as phase of the slow oscillation at the time of stimulation (e.g., Batterink et al, 2016; Ong et al,

2016) becomes difficult with more complex sounds. The length of the stimuli, possibly including a fade in and fade out to avoid arousal and its detrimental effects on the efficacy of stimulation (Göldi & Rasch, 2019), makes it difficult to pinpoint when the stimulation would be with regards to the phase of the slow oscillation. These shortcomings do not prevent the data from being useful for other analysis tools: Data-driven techniques, such as empirical mode decomposition (e.g., Stallone et al, 2020), appear to show promise, allowing broader applicability of these new analysis methods to sleep data.

In general, running a complex study like the "nap study" teaches many valuable lessons about research at a more fundamental level. It emphasised the value, and beauty, of interdisciplinary scientific approaches and creativity: for this study, musicians and circadian researchers provided valuable feedback and contributed considerations that are often overlooked in their implications on experimental work. An example of this is the dependence of the cognitive abilities targeted by the experimental tasks on the quantity, quality, and recency of the sleep the participant had. Further, building up a sleep laboratory requires creativity, flexibility, and co-operation: Sleep studies require a different infrastructure to other cognitive studies, presenting with various sources of challenges. It needs creativity to incorporate a bed or bed-like structure into existing testing spaces, as well as the willingness of fellow researchers to support such restructuring. Different equipment challenges might include difficulties in sending EEG triggers through longer cables or finding headphones that aren't interfering with sleep or the measured EEG signal. Further, and importantly, sleep studies require a lot of co-operation of fellow researchers to keep noise levels in the testing areas low enough to facilitate a participant falling and remaining asleep. This leads on to another lesson that sleep studies stressed: The human aspect of experiments is often not

considered beyond clinical settings and patient studies. Within the "nap study" the participants' well-being, comfort, and interest in the study was of great importance. Interest was key because the testing duration of a participant alone took an entire day, emphasising the need for the data to be usable and the participants returning for the second session. However, more important was the comfort of the participant: Participants had to feel safe and comfortable enough to fall asleep in an unfamiliar location, with the discomfort of an EEG cap and headset on, and knowing that the experimenter was sitting just outside monitoring them. Building a relationship of trust and ease with the participants was crucial to enable such levels of comfort.

To summarise, the results of the study on selective reactivation during consolidation are congruent with observations and findings in the literature on the intricate interplay of attention and memory. Further, the present study supports the potential of nap studies for the study of sleep and sleep-related effects. However, some of the current findings clearly set the study apart from many commonly held beliefs in the field: Sleep did not universally help the consolidation of episodic memories, and sleep-based stimulation was not purely beneficial. Indeed, prolonged stimulation may even be detrimental for the flexible use of memories in guiding attention. The results do not preclude a functional role of sleep-based consolidation for memory and do suggest that memories might be malleable for "attentional" input during the consolidation process. However, they suggest that the relationship might not be as straightforward as perhaps assumed, and that certain processes during memory consolidation might even render detrimental effects on the functional utility of the consolidated memories.

Retrocues during memory retrieval

After targeting encoding and consolidation processing in previous studies, the third set of studies attempted to manipulate how attention was deployed during retrieval. The aim was to illuminate attentional effects on the process of retrieval and the retrieved memories. Pertaining to the retrieved memories, there appears to be a conceptual distinction: On the one hand, the effects of an attentional bias at retrieval could modulated the recalled mnemonic information in response to the attentional bias. And on the other hand, or additionally, the attentional bias during retrieval could have a prolonged influence on the content of the corresponding memory. The discussed set of studies addressed these questions, by using retrocues to bias attention for the selective retrieval of long-term memories.

First and foremost, the current set of results confirms the efficacy of retrocues in long-term memory: the selective retrieval of long-term memories was successfully biased by guiding internal attention to goal-relevant targets. The robust and substantial benefit of informative retrocues was replicated across all three studies, allowing faster access to the relevant mnemonic information. The predominant outcome measure of the retrocueing benefit was reaction time, with no consistent effects observed for accuracy measures. This was unsurprising given the emphasis on speeded reaction times, and the layout of the task: The retrocues were of a spatial nature, indicating the remembered hemisphere of the target object, before an identity judgement was made about which object had been present in the scene. This perceptual decision was very simple, leading to ceiling accuracy performance values across conditions. Different task designs could emphasise the accuracy component by increasing the difficulty of the perceptual decision. This could be achieved, for example, by forcing a decision between similar items (examples of the same category or different viewpoints of the

same item) or decreasing visibility of the items (lowering saturation or masking). A different option to emphasise accuracy measures targets the retrocue itself: varying the reliability of the retrocue increases the task difficulty, which in turn should emphasise and cause a wider distribution of accuracy data.

The latter approach also relates to a different question about the functionality of retrocues: so far, the efficacy of a retrocue was discussed in reference to a benefit for the memory, yet could there also be a detrimental effect? Study 3 raises this issue by contrasting informative retrocues that are either accurately predicting the upcoming target selection or are misleading. The impact of the retrocue appeared amplified in the contrast of accurate and misleading information, when compared to the benefit of informative relative to neutral retrocues used in Study 1 and Study 2. The results indicated that the internal attentional guidance during retrieval might also lead to a cost in performance, if the anticipated events do not come to pass as expected. If we are distracted or use erroneous information while we attempt to retrieve particular (aspects of) memories, the retrieved memory might not help us to achieve our task or might even distract us further by tempting us to "go down memory lane". Indeed, one could argue that the attentional biases towards negative information in depression might cause the selective retrieval of negative memories, a focus on which in turn might contribute to the disrupted cognitive processes in depression (e.g., see Hammar & Årdal, 2009). In working memory, the cost of invalid retrocues has already been shown (Astle et al, 2012). Within long-term memory, proof of potentially detrimental effects of orienting attention during retrieval still requires further study. In the current study using misleading retrocues, it is difficult to disentangle the putatively detrimental effects of the misleading retrocues from beneficial effects of the accurate retrocues. A natural follow-up study would contrast accurate, misleading, and neutral retrocues within one

task to allow a direct comparison of the beneficial and detrimental effects of the informative cues relative to the neutral cues. As the inclusion of misleading, as opposed to neutral, retrocues led to significant retrocueing effects on the speed and accuracy of identifying the memory target, as well as on the precision of positioning the target in context, I would predict a likely cost of misleading retrocues relative to neutral ones.

The reflection on retrocueing effects on memory precision links to the second topic of consideration: Beyond affecting access to the memory, can internally oriented attention change the quality of the retrieved information? The results of Study 3 showed that the nature of the retrocue indeed had an impact on the spatial precision with which the target object was positioned in the scene context. This finding parallels observations in working memory studies on retrocues (e.g., Gunseli et al, 2015; Pertzov et al, 2017; see also Zokaei et al, 2014; but cf. Myers et al, 2017), yet extends our understanding of the impact of attention during retrieval further: Even the contents of stable, long-term memories can be biased by attentional selection during retrieval. It is noteworthy that this precision effect was not a necessary consequence of the successful utilisation of the retrocue: Higher precision in response to accurate information was only observed in contrast to misleading information (Study 3), not in contrast to neutral retrocues (Study 2). This observation warrants further discussion.

It is possible that the retrocueing effects on precision only became visible through an accumulation of benefits from valid retrocues and costs from misleading retrocues. Consequently, it is possible that limited sensitivity of the precision measure contributed to the absence of a significant effect in Study 2. The current task used a straightforward spatial recall task and resulting average Euclidean distance to assess retrocueing effects on memory quality. Average Euclidean distance is susceptible to disproportionate influences of a few guesses or swap errors between the two target locations. A more

sensitive measure, for example modelling approaches taking swap errors and guesses into account, might allow a more appropriate assessment of memory and might be able to test with greater sensitivity for improvements in spatial precision resulting only from benefits of valid spatial retrocues relative to neutral retrocues.

It should be noted that despite the current retrocues being spatial in nature, there is nothing that precludes retrocueing different features or dimensions, given the right task design. Colour, higher-order concepts (living or non-living), or categories (e.g., types of objects, animal) could all be feasible cues. Indeed, many of these measures lend themselves to sensitive precision measures, such as a colour wheel, orientation measures, or angle of point of view (see Richter et al, 2016, for examples of applicable continuous precision measures). There are also ways of creating precision measures if continuous scales are not an option: Distinctions could be drawn between examples of a category with varying perceptual similarity, or items morphed to different degrees, obscuring conceptual or categorical identity.

The effects of the retrocue on the accuracy of the retrieved memory can also be interpreted in regard to the precision of the retrocue itself. In this task, the retrocue was beneficial when it directed attention to the hemisphere of the appearing target ("valid") vs. when pointing away from it ("invalid"). Still to be determined is whether the resulting effect was categorical or continuous. Is the effect of the retrocue dependent on the precision of the retrocue? Is a spatial retrocue more effective if it points directly to the remembered location than if it highlights a larger patch within the scene or an entire side? Experimentally, this hypothesis could easily be assessed by varying the fit of the retrocue to the target locations. The size of the retrocueing effect on memory accuracy, and indeed the access to the memory, might vary as a function of the specificity of the retrocue. The inverse of this thought process concerns the specificity

of the output of the retrocue: Is the attentional highlight the retrocue appears to afford restricted to the targeted feature? This would limit spatial retrocues to enable spatial locations while colour-based retrocues could only bias colour perception. Or can a retrocue highlight complex, multifaceted memories as a whole that might encompass multiple features? This consideration is especially important for episodic memories, consisting frequently of a conglomerate of spatial, temporal, and contextual information. The pertinent hypothesis would be that a spatial retrocue would be able to bias selective retrieval to associated other features of that episodic memory, such as colour, emotion, time, etc.

One dimension that may play an important role in determining the types of benefits and/or costs conferred by retrocues is the task demands. Studies in the domain of perceptual attention highlight how different task requirements influence the types of behavioural effects observed. For example, tasks requiring speeded responses to easily identifiable or discriminable stimuli promote changes in response times. In contrast, tasks requiring difficult perceptual identification or discrimination under non-speeded conditions promote changes to perceptual sensitivity. The studies in Chapter IV all employed the same task of retrieving a salient item from a contextual memory under speeded condition, thus emphasising response speed and making reaction times the most sensitive measure. In future versions of the task, it would be possible to manipulate the mnemonic and response demands systematically. An increased emphasis on perceptual sensitivity could be achieved, for example, by presenting the target objects within the scene context but not necessarily the exact learnt location. Participants would have to report whether the object was in the learnt location or not.

Is it feasible to suppose that the selectively retrieved content of the memory could be biased, or even altered, to match the anticipated needs? Indeed, within working

memory, it has been noted that the nature of prioritisation of items depended on the intended use (see Nobre & Stokes, 2019), for example, by distorting remembered colour to optimise behaviour (Yu & Geng, 2019). Further, long-term memories have been shown to be malleable by perceived task relevance during the encoding and consolidation stage: Features attended to during encoding show selective enhancement (Uncapher & Rugg, 2009; Uncapher et al 2011) and consolidation is biased towards memories based on perceived future relevance (Diekelman & Born, 2010) or reward (Oudiette et al, 2013). Does this task-specific modulation also occur during retrieval? Can selectively attending to and retrieving aspects of a stable long-term memory modulate its content to match behavioural needs as best possible? Famously, the relationship of specific retrieval questions and modulation of the resulting memory has direct consequences for eyewitness memories (e.g., Loftus, 2005). In their classic car crash experiment, Loftus and Palmer (1974) found significant differences in the content of the retrieved memory depending on how participants were asked about circumstances surrounding a car crash. The implications of these observations for researchers, clinicians, and officials alike are that one needs to be careful about the focus of internal attention during retrieval, and the nature and information in the cues used to orient such focus, as this could impact the content of the retrieved memory.

This leads on to the third aspect of discussing the results: If selective attention guided by retrocues can have an impact on the immediately retrieved mnemonic content, do these effects also have lasting consequences for the long-term memory? The present set of results did not give a clear-cut answer to this question. It has to be noted that any consequences of the retrocue were not as reliable or substantial as the effects discussed in the previous sections. I observed some benefits of probing during the retrocueing task on the speed of later retrieval (Study 1), a spatial precision benefit

if the retrocue was informative (Study 2), and an accuracy effect for later identification depending on which object the retrocue was predicting to appear (Study 3). None of the observed effects were reliable over all studies, and potential reasons are considered in the discussion of Chapter IV. What can we deduce from these diverse results? Firstly, unless all of these results were due to false positives, selectively retrieving a memory can have prolonged impact on several aspects of the memory. Secondly, potential effects appeared weak and not a determinative consequence of the selective retrieval. How can these results be reconciled with the substantial effects of selective attention on resulting memory observed in eyewitness studies (see Loftus, 2005)? Multiple aspects could contribute to the differences in observation. Relating back to the previously discussed role of task specificity, it is possible that biasing information could be integrated to a higher degree if it is perceived to be important for the memory. The leading questions in eyewitness studies have direct relevance for the interpretation of the memory, and consequently the selectively attended information could be integrated more lastingly. Further, there might be a cumulative effect of repetition. Recalling a memory multiple times with a certain attentional focus might increase distortion for this particular aspect over time.

In a nutshell, the study on selective attention during retrieval (Chapter IV) gave the first and convincing evidence for the efficacy of retrocues to investigate selective retrieval within long-term memory. This demonstrates that memory retrieval is susceptible to attentional modulation, with immediate effects and potentially longer standing impacts on the content of the memory. The set of studies therefore represents a great starting point for investigating the attentional impact on selective memory retrieval. As discussed in the chapter, the new experimental approach shows great promise for integrating different lines of research and various methodologies, ranging

from eyetracking to fMRI. This versatility makes the set-up appropriate for a multitude of conceptual, behavioural, and neural enquiries. The engaging nature of the tasks, especially the learning task, also makes it adaptable for different kinds of studies and populations, such as children or patients.

The way forward

Taken together, the body of my doctoral work highlighted that our memories are indeed subject to attentional inputs during each of the essential stages of encoding, consolidation and retrieval. Interestingly, each interference resulted in some prolonged effects on the memory, even if these effects were too subtle to be picked up by coarse behavioural measures. Nonetheless, the implications of this are striking: Our memories appear to be shaped by the attentional input they receive. What, if any, limitations does this putative mnemonic shaping have? Given enough attentional input under the right circumstances, can we change memory content fluently? Even if we could manipulate merely the context or interpretation of a memory via these attentional mechanisms, the implications for treatment could be enormous. Complementing the role of memory and attention interactions in the treatment of PTSD or depression as discussed in the introduction, shaping of memory would be an incredible tool to place traumatic memories in a different light. Attentional redirection might allow strengthening different interpretations of the memory or redirecting feelings of, for example, shame and guilt, often associated with trauma memories. For this aim, further research is required to examine 1) whether attentional direction in memory can be amplified for example by repeated exposure, 2) whether such attentional modulations produce lasting changes in the memory, both functionally and structurally, and 3) how these observations could be translated into applied tools for treatment.

A second observation derived from the three strands of research discussed in this doctoral thesis concerns the use of the memory itself: It appears, perhaps unsurprisingly, that not the veracity of the memory is decisive but its utility. Are our memories indeed changed and adapted proactively throughout their establishment to meet presumed future requirements? In reverse, are attentional modulations limited to the current use of the memory, if their utility is considered limited to that instance? These ideas are congruent with the sleep literature, highlighting relevant memories for consolidation. Such observations have great implications for the natural development of long-term memories across our lifespan. Further research is needed to address 1) if perceived future utility amplifies or defines lasting modulations to the long-term memory, 2) at which stages of memory establishment such proactive modulations are most efficacious, and 3) to what extent presumed proactive adaptation of memories can lastingly alter memories across time.

Concluding remarks

The foundational motivation for this doctoral thesis was to illuminate the highly interactive and intricate relationship of memory and attention. In particular I sought to explore whether attentional impacts could influence the establishment of memory across encoding, consolidation, and retrieval, wondering about potential lasting moderation of the quality and functional use of the memory. Indeed, I found that attentional modulation had some lasting effects throughout all three stages: Distraction biased memory encoding, distorted elements of the resulting memory, and required divergent processing in its use as attentional guide. Targeted memory reactivation during consolidation was associated with protection against forgetting but loss of mnemonic function in guiding attention. Retrocues enabled and supported the selective

retrieval of task-relevant memoranda, sharpening the mnemonic content to fit behavioural requirements. In contrast, I also observed limitations in our ability to suppress distraction and witnessed great resilience in the encoding processes in the face of such distraction. The influence of sleep and sleep-based stimulation on consolidation was not replicated and it was suggested that effects may not necessarily always be beneficial to the memory. And whilst selective retrieval might allow improved retrieval of elements of particular interest, this bias was not conclusively shown to influence the nature of the existing memory. Altogether the body of research in this thesis emphasises that memory and attention do indeed not exist separately but work in concert. The findings contribute to our understanding of how memory and attention interact, highlighting the multifaced nature of these mutual influences, and open doors for exciting new avenues of research.

References

Abbot, E. (1909). On the analysis of the factors of recall in the learning process. *Psychological Monographs: General and Applied, 11*(1), 159-177.

Aly, M., & Turk-Browne, N. B. (2016a). Attention promotes episodic encoding by stabilizing hippocampal representations. *PNAS*, E420-429.

Aly, M., & Turk-Browne, N. B. (2016b). Attention Stabilizes Representations in the Human Hippocampus. *Cerebral Cortex*, *26*(February), 783–796.

Aly, M., & Turk-Browne, N. B. (2017). How Hippocampal Memory Shapes, and Is shaped by, Attention. In D. E. Hannula & M. C. Duff (Eds.), *The Hippocampus from Cells to Systems.* (pp. 369–403). Switzerland: Springer International Publishing AG.

America Electroencephalographic Society (1991). American electroencephalographic society guidelines for standard electrode position nomenclature, *Journal of Clinical Neurophysiology, 8*, 200–202.

American Psychiatric Association (2000). *Diagnostic and statistical manual of mental disorders (4th ed. TR)*. Washington, DC: APA.

Anderson, B. A. (2016). The attention habit: how reward learning shapes attentional selection. *Annals of the New York Academy of Sciences*, *1369*, 24–39.

Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering Can Cause Forgetting: Retrieval Dynamics in Long-Term Memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(5), 1063–1087.

Anderson, M. C., & Hanslmayr, S. (2014). Neural mechanisms of motivated forgetting. *Trends in Cognitive Sciences*, *18*(6), 279–292.

Antony, J. W., Gobel, E. W., O'Hare, J. K., Reber, P. J., & Paller, K. A. (2012). Cued memory reactivation during sleep influences skill learning. *Nature Neuroscience*, *15*(8), 1114–1116.

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*, 1736–1743.

Antony, J. W., Schönauer, M., Staresina, B. P., & Cairney, S. A. (2019). Sleep Spindles and Memory Reprocessing. *Trends in Neurosciences*, *42*(1), 1–3.

Arita, J. T., Carlisle, N. B., & Woodman, G. F. (2012). Templates for rejection: Configuring attention to ignore task-irrelevant features. *Journal of Experimental Psychology. Human Perception and Performance*, *38*(3), 1–10.

Astle, D. E., Summerfield, J., Griffin, I., & Nobre, A. C. (2012). Orienting attention to locations in mental representations. *Attention, Perception, and Psychophysics*, *74*(1), 146–162.

Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. *Psychology of learning and motivation*, *2*(4), 89-195.

Averbach, E., & Coriell, A.S. (1961). Short-term memory in vision. *Bell System Technical Journal, 40*, 309–328.

Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*(8), 437–443.

Backer, K. C., & Alain, C. (2012). Orienting Attention to Sound Object Representations Attenuates Change Deafness. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(6), 1554.

Backer, K. C., Binns, M. A., & Alain, C. (2015). Neural dynamics underlying attentional orienting to auditory representations in short-term memory. *Journal of Neuroscience, 35*, 1307–1318.

Baird, B., Mota-Rolim, S. A., & Dresler, M. (2019). The cognitive neuroscience of lucid dreaming. *Neuroscience and Biobehavioral Reviews*, *100*, 305–323.

Bar, M. (2009). The proactive brain: memory for predictions. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *364*, 1235–1243.

Batterink, L. J., Creery, J. D., & Paller, K. A. (2016). Phase of Spontaneous Slow Oscillations during Sleep Influences Memory-Related Processing of Auditory Cues. *The Journal of Neuroscience*, *36*(4), 1401–1409.

Bäuml, K. H. T. (2019). Context Retrieval as a Critical Component in Selective Memory Retrieval. *Current Directions in Psychological Science*, *28*(2), 177–182.

Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9(10)*:7, 1–11.

Beck, V. M., & Hollingworth, A. (2015). Evidence for negative feature guidance in visual search is explained by spatial recoding. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(5), 1190–1196.

Beck, V. M., Luck, S. J., & Hollingworth, A. (2018). Whatever you do, don't look at the... Evaluating guidance by an exclusionary attentional template. *Journal of Experimental Psychology: Human Perception and Performance*, *44*(4), 645–662.

Bendor, D., & Wilson, M. A. (2012). Biasing the content of hippocampal replay during sleep. *Nature Neuroscience*, *15*(10), 1439–1444.

Biederman, I. (1972). Perceiving real-world scenes. *Science*, *177*(4043), 77-80.

Bomyea, J., Johnson, A., & Lang, A. J. (2017). Information Processing in PTSD: Evidence for Biased Attentional, Interpretation, and Memory Processes. *Psychopathology Review*, *4*(3), 218–243.

Born, J., Rasch, B., & Gais, S. (2006). Sleep to Remember. *The Neuroscientist*, *12*(5), 410–424.

Born, J., & Wilhelm, I. (2012). System consolidation of memory during sleep. *Psychological Research*, *76*, 192–203.

Broadbent, D. (1958). *Perception and Communication*. London: Pergamon Press.

Brockmole, J. R., Castelhano, M. S., & Henderson, J. M. (2006). Contextual cueing in naturalistic scenes: Global and local contexts. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *32*(4), 699–706.

Buffalo, E. A. (2015). Bridging the Gap Between Spatial and Mnemonic Views of the Hippocampal Formation. *Hippocampus*, *25*, 713–718.

Burgess, N., Maguire, E. A., & O'Keefe, J. (2002). The human hippocampus and spatial and episodic memory. *Neuron*, *35*(4), 625–641.

Byrne, J. H. (2017). *Learning and memory: a comprehensive reference*. Academic Press.

Cabeza, R. (2008). Role of parietal regions in episodic memory retrieval: The dual attentional processes hypothesis. *Neuropsychologia*, *46*(7), 1813–1827.

Cabeza, R., Ciaramelli, E., Olson, I. R., & Moscovitch, M. (2008). The parietal cortex and episodic memory: an attentional account. *Nature Reviews Neuroscience*, *9*, 613–625.

Cabeza, R., & Nyberg, L. (2000). Imaging cognition II: An empirical review of 275 PET and fMRI studies. *Journal of cognitive neuroscience*, *12*(1), 1-47.

Carlisle, N. B., & Woodman, G. F. (2011). Automatic and strategic effects in the guidance of attention by working memory representations. *Acta Psychologica*, *137*(2), 217–225.

Carrier, M., & Pashler, H. (1992). The influence of retrieval on retention. *Memory & Cognition*, *20*(6), 633–642.

Cash, S. S., Halgren, E., Dehghani, N., Rossetti, A. O., Thesen, T., Wang, C., Devinsky, O., Kuzniecky, R., Doyle, W., Madsen, J. R., Bromfield, E., Erőss, L., Halász, P., Karmos, G., Csercsa, R., Wittner, L., & Ulbert, I. (2009). The human K-complex represents an isolated cortical down-state. *Science*, *324*(5930), 1084–1087.

Chelazzi, L., Marini, F., Pascucci, D., & Turatto, M. (2019). Getting rid of visual distractors: the why, when, how, and where. *Current Opinion in Psychology*, *29*, 135–147.

Chelazzi, L., Perlato, A., Santandrea, E., & Della Libera, C. (2013). Rewards teach visual selective attention. *Vision Research*, *85*, 58–72.

Chen, Z., Wimmer, R. D., Wilson, M. A., & Halassa, M. M. (2016). Thalamic Circuit Mechanisms Link Sensory Processing in Sleep and Attention. *Frontiers in Neural Circuits*, *9*(1), 1–10.

Cherry, C. (1953). Some experiments on the recognition of speech, with one and with two ears. *The Journal of the Acoustical Society of America, 25,* 975–979.

Chun, M. M. (2000). Contextual cueing of visual attention. *Trends in Cognitive Sciences*, *4*(5), 170–178.

Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology*, *62*, 73–101.

Chun, M. M., & Jiang, Y. (1998). Contextual cueing: implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, *36*(1), 28–71.

Chun, M. M., & Johnson, M. K. (2011). Memory: Enduring traces of perceptual and reflective attention. *Neuron*, *72*(4), 520–535.

Chun, M. M., & Phelps, E. A. (1999). Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. *Nature Neuroscience*, *2*(9), 844–847.

Chun, M. M., & Turk-Browne, N. B. (2007). Interactions between attention and memory. *Current opinion in neurobiology*, *17*(2), 177-184.

Ciaramelli, E., Grady, C. L., & Moscovitch, M. (2008). Top-down and bottom-up attention to memory: A hypothesis (AtoM) on the role of the posterior parietal cortex in memory retrieval. *Neuropsychologia*, *46*(7), 1828–1851.

Clayton, N. S., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, *395*(6699), 272-274.

Cohen, N. J., Eichenbaum, H. (1993). *Memory, amnesia, and the hippocampal system.* MIT press.

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews. Neuroscience*, *3*(March), 201–215.

Corkin, S. (2002). What's new with the amnesic patient H.M.? *Nature Reviews Neuroscience*, *3*(2), 153–160.

Corkin, S. (2013). *Permanent present tense: The man with no memory, and what he taught the world*. Penguin UK.

Craik, F. I. M., Govoni, R., Naveh-Benjamin, M., & Anderson, N. D. (1996). The Effects of Divided Attention on Encoding and Retrieval Processes in Human Memory. *Journal of Experimental Psychology*, *125*(2), 159–180.

Creery, J. D., Oudiette, D., Antony, J. W., & Paller, K. A. (2015). Targeted Memory Reactivation during Sleep Depends on Prior Learning. *Sleep*, *38*(5), 755–763.

Crowley, R., Bendor, D., & Javadi, A.-H. (2019). A review of neurobiological factors underlying the selective enhancement of memory at encoding, consolidation, and retrieval. *Progress in Neurobiology*, *179*(101615), 1–12.

Cunningham, C. A., & Egeth, H. E. (2016). Taming the White Bear: Initial Costs and Eventual Benefits of Distractor Inhibition. *Psychological Science*, *27*(4), 476–485.

Curran, T., & Dien, J. (2003). Differentiating amodal familiarity from modality-specific memory processes: An ERP study. *Psychophysiology*, *40*(6), 979–988.

Daffner, K. R., Zhuravleva, T. Y., Sun, X., Tarbi, E. C., Haring, A. E., Rentz, D. M., & Holcomb, P. J. (2012). Does modulation of selective attention to features reflect enhancement or suppression of neural activity? *Biological Psychology*, *89*(2), 398–407.

Dalgleish, T., & Watts, F. N. (1990). Biases of attention and memory in disorders of anxiety and depression. *Clinical Psychology Review*, *10*, 589–604.

Danker, J. F., & Anderson, J. R. (2010). The ghosts of brain states past: Remembering reactivates the brain regions engaged during encoding. *Psychological Bulletin*, *136*(1), 87–102.

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.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*, 193–222.

Diekelmann, S., & Born, J. (2010). The memory function of sleep. *Nature Reviews*, *11*(2), 114–126.

Doallo, S., Patai, E. Z., & Nobre, A. C. (2013). Reward Associations Magnify Memory-based Biases on Perception. *Journal of Cognitive Neuroscience*, *25*(2), 245–257.

Doherty, B. R., Patai, E. Z., Duta, M., Nobre, A. C., & Scerif, G. (2017). The functional consequences of social distraction: Attention and memory for complex scenes. *Cognition*, *158*, 215–223.

Donohue, S. E., Bartsch, M. V, Heinze, H.-J., Schoenfeld, M. A., & Hopf, J.-M. (2018). Cortical Mechanisms of Prioritizing Selection for Rejection in Visual Search. *The Journal of Neuroscience*, *38*(20), 4738–4748.

Drosopoulos, S., Schulze, C., Fischer, S., & Born, J. (2007). Sleep's function in the spontaneous recovery and consolidation of memories. *Journal of Experimental Psychology: General*, *136*(2), 169–183.

Drosopoulos, S., Wagner, U., & Born, J. (2005). Sleep enhances explicit recollection in recognition memory. *Learning & Memory*, *12*(1), 44–51.

Dudai, Y. (2004). The Neurobiology of Consolidations, Or, How stable is the engram? *Annual Review of Psychology*, *55*, 51–86.

Dudukovic, N. M., & Wagner, A. D. (2006). Attending to Remember and Remembering to Attend. *Neuron*, *49*(6), 784–787.

Duncan, K., Ketz, N., Inati, S. J., & Davachi, L. (2012). Evidence for Area CA1 as a Match/Mismatch Detector: A High-Resolution fMRI Study of the Human Hippocampus. *Hippocampus*, *22*, 389–398.

Ebbinghaus, H. (1885). *Über das Gedächtnis: Untersuchungen zur experimentellen Psychologie*. Duncker & Humblot.

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

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*, 716–722.

Falkenstein, M., Hoormann, J., & Hohnsbein, J. (1999). ERP components in Go/Nogo tasks and their relation to inhibition. *Acta Psychologica*, *101*, 267–291.

Fan, J. E., & Turk-Browne, N. B. (2015). Incidental Biasing of Attention From Visual Long-Term Memory. *Journal of Experimental Psychology*, *41*(6), 1–8.

Federmeier, K. D., Mai, H., & Kutas, M. (2005). Both sides get the point: Hemispheric sensitivities to sentential constraint. *Memory and Cognition*, *33*(5), 871–886.

Fernandes, M. A., & Moscovitch, M. (2000). Divided Attention and Memory: Evidence of Substantial Interference Effects at Retrieval and Encoding. *Journal of Experimental Psychology*, *129*(2), 155–176.

Ferrante, O., Patacca, A., Di Caro, V., Della Libera, C., Santandrea, E., & Chelazzi, L. (2018). Altering spatial priority maps via statistical learning of target selection and distractor filtering. *Cortex*, *102*, 67–95.

Ferree, T. C. (2000). Spline interpolation of the scalp EEG. *Secondary Titl. EGI*. http://www.egi.com/Technotes/SplineInterpolation.pdf

Fischer, S., & Born, J. (2009). Anticipated Reward Enhances Offline Learning During Sleep. *Journal of Experimental Psychology*, *35*(6), 1586–1593.

Fischer, M., Moscovitch, M., & Alain, C. (2021). A systematic review and meta-analysis of memory-guided attention: Frontal and parietal activation suggests involvement of fronto-parietal networks. *Wiley Interdisciplinary Reviews: Cognitive Science*, *12*(1), e1546:1-25.

Foa, E. B., & Rothbaum, B. O. (2001). *Treating the trauma of rape: Cognitive-behavioral therapy for PTSD*. Guilford Press.

Fongaro, E., & Rose, J. (2020). Crows control working memory before and after stimulus encoding. *Scientific Reports*, *10*(3253), 1–10.

Foxe, J. J., & Snyder, A. C. (2011). The role of alpha-band brain oscillations as a sensory suppression mechanism during selective attention. *Frontiers in Neuroscience*, *2*, 154: 1-13.

Fukuda, K., & Woodman, G. F. (2017). Visual working memory buffers information retrieved from visual long-term memory. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(20), 5306–5311.

Gaspar, J. M., & McDonald, J. J. (2014). Suppression of Salient Objects Prevents Distraction in Visual Search. *The Journal of Neuroscience*, *34*(16), 5658–5666.

Gaspelin, N., Leonard, C. J., & Luck, S. J. (2015). Direct Evidence for Active Suppression of Salient-but-Irrelevant Sensory Inputs. *Psychological Science*, *26*(11), 1740–1750.

Gaspelin, N., Leonard, C. J., & Luck, S. J. (2017). Suppression of overt attentional capture by salient-but-irrelevant color singletons. *Attention, Perception, & Psychophysics*, *79*, 45–62.

Gaspelin, N., & Luck, S. J. (2018a). Electrophysiological and behavioral evidence of suppression of salient-but-irrelevant stimuli. *Journal of Cognitive Neuroscience*, *30*(9), 1265-1280.

Gaspelin, N., & Luck, S. J. (2018b). The Role of Inhibition in Avoiding Distraction by Salient Stimuli. *Trends in Cognitive Sciences*, *22*(1), 79–92.

Gaspelin, N., & Luck, S. J. (2019). Inhibition as a potential resolution to the attentional capture debate. *Current Opinion in Psychology*, *29*, 12–18.

Gazzaley, A., & Nobre, A. C. (2012). Top-down modulation: Bridging selective attention and working memory. *Trends in Cognitive Sciences*, *16*(2), 129–135.

Girardeau, G., Benchenane, K., Wiener, S. I., Buzsáki, G., & Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, *12*(10), 1222–1223.

Goldfarb, E. V, Chun, M. M., & Phelps, E. A. (2016). Memory-Guided Attention: Independent Contributions of the Hippocampus and Striatum. *Neuron*, *89*, 1–8.

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), 1–7.

Gotlib, I. H., & Joormann, J. (2010). Cognition and Depression: Current Status and Future Directions. *Annual Review of Clinical Psychology*, *6*, 285–312.

Gotlib, I. H., Krasnoperova, E., Neubauer Yue, D., & Joormann, J. (2004). Attentional Biases for Negative Interpersonal Stimuli in Clinical Depression. *Journal of Abnormal Psychology*, *113*(1), 127–135.

Graf, P., & Schacter, D. L. (1985). Implicit and Explicit Memory for New Associations in Normal and Amnesic Subjects. *Journal of Experimental Psychology*, *11*(3), 501–518.

Greving, S., & Richter, T. (2018). Examining the Testing Effect in University Teaching: Retrievability and Question Format Matter. *Frontiers in Psychology*, *9*, 1–10.

Griessenberger, H., Hoedlmoser, K., Heib, D. P. J., Lechinger, J., Klimesch, W., & Schabus, M. (2012). Consolidation of temporal order in episodic memories. *Biological Psychology*, *91*(1), 150–155.

Griffin, I. C., & Nobre, A. C. (2003). Orienting Attention to Locations in Internal Representations. *Journal of Cognitive Neuroscience*, *15*(8), 1176–1194.

Günseli, E., & Aly, M. (2020). Preparation for upcoming attentional states in the hippocampus and medial prefrontal cortex. *ELife*, (9), e53191:1-39.

Gunseli, E., van Moorselaar, D., Meeter, M., & Olivers, C. N. L. (2015). The reliability of retro-cues determines the fate of noncued visual working memory representations. *Psychonomic Bulletin & Review*, *22*, 1334–1341.

Halassa, M. M., Chen, Z., Wimmer, R. D., Brunetti, P. M., Zhao, S., Zikopoulos, B., Wang, F., Brown, E. N., & Wilson, M. A. (2014). State-Dependent Architecture of Thalamic Reticular Subnetworks. *Cell*, *158*(4), 808–821.

Hammar, Å., & Årdal, G. (2009). Cognitive functioning in major depression – a summary. *Frontiers in Human Neuroscience*, *3*, 26: 1-7.

Hannula, D. E., Althoff, R. R., Warren, D. E., Riggs, L., & Cohen, N. J. (2010). Worth a glance: using eye movements to investigate the cognitive neuroscience of memory. *Frontiers in Human Neuroscience*, *4*(October), 1–16.

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

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

Harper, J., Malone, S. M., & Bernat, E. M. (2014). Theta and delta band activity explain N2 and P3 ERP component activity in a go/no-go task. *Clinical Neurophysiology*, *125*(1), 124–132.

Hashimoto, R., Abe, N., Ueno, A., Fujii, T., Takahashi, S., & Mori, E. (2012). Changing the Criteria for Old / New Recognition Judgments Can Modulate Activity in the Anterior Hippocampus. *Hippocampus*, *22*, 141–148.

Hazy, T. E., Frank, M. J., & O'Reilly, R. C. (2007). Towards an executive without a homunculus: Computational models of the prefrontal cortex/basal ganglia system. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1485), 1601–1613.

Hickey, C., Di Lollo, V., & McDonald, J. J. (2009). Electrophysiological indices of target and distractor processing in visual search. *Journal of Cognitive Neuroscience*, *21*(4), 760–775.

Higo, T., Mars, R. B., Boorman, E. D., Buch, E. R., & Rushworth, M. F. S. (2011). Distributed and causal influence of frontal operculum in task control. *PNAS*, *108*(10), 4230–4235.

Hillyard, S. A., Vogel, E. K., & Luck, S. J. (1998). Sensory gain control (amplification) as a mechanism of selective attention: Electrophysiological and neuroimaging

evidence. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *353*(1373), 1257–1270.

Hollingworth, A. (2006). Scene and Position Specificity in Visual Memory for Objects. *Journal of Experimental Psychology*, *32*(1), 58–69.

Hollingworth, A. (2009). Two forms of scene memory guide visual search: Memory for scene context and memory for the binding of target object to scene location. *Visual Cognition*, *17*(1–2), 273–291.

Holmes, E. A., James, E. L., Coode-Bate, T., & Deeprose, C. (2009). Can Playing the Computer Game '"Tetris"' Reduce the Build-Up of Flashbacks for Trauma? A Proposal from Cognitive Science. *PLoS ONE*, *4*(1), 1–6.

Horne, J. A., & Östberg, O. (1976). A self-assessment questionnaire to determine morningness-eveningness in human circadian rhythms. *International Journal of Chronobiology, 4*, 97-100.

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.

Iber, C., Ancoli-Israel, S., Chesson, A. L., & Quan, S. F. (2007). *The AASM Manual for the Scoring of Sleep and Associated Events*. Westchester, Ill: American Academy of Sleep Medicine.

Inostroza, M., & Born, J. (2013). Sleep for Preserving and Transforming Episodic Memory. *Annual Review of Neuroscience*, *36*, 79–102.

Ipata, A. E., Gee, A. L., Gottlieb, J., Bisley, J. W., & Goldberg, M. E. (2006). LIP responses to a popout stimulus are reduced if it is overtly ignored. *Nature Neuroscience*, *9*(8), 1071–1076.

Iyadurai, L., Blackwell, S. E., Meiser-Stedman, R., Watson, P. C., Bonsall, M. B., Geddes, J. R., Nobre, A. C., & Holmes, E. A. (2018). Preventing intrusive memories after trauma via a brief intervention involving Tetris computer game play in the emergency department: a proof-of-concept randomized controlled trial. *Molecular Psychiatry*, *23*(3), 674–682.

James W. 1890. *The Principles of Psychology*. New York: Holt.

James, E. L., Bonsall, M. B., Hoppitt, L., Tunbridge, E. M., Geddes, J. R., Milton, A. L., & Holmes, E. A. (2015). Computer Game Play Reduces Intrusive Memories of Experimental Trauma via Reconsolidation-Update Mechanisms. *Psychological Science*, *26*(8), 1201–1215.

Jenkins, J. G., & Dallenbach, K. M. (1924). Obliviscence during Sleep and Waking. *The American Journal of Psychology*, *35*(4), 605–612.

Jensen, O., & Mazaheri, A. (2010). Shaping Functional Architecture by Oscillatory Alpha Activity: Gating by Inhibition. *Frontiers in Human Neuroscience*, *4*, 186:1–8.

Jiang, Y. V. (2018). Habitual versus goal-driven attention. *Cortex*, *102*, 107–120.

Johnson, J. A., & Zatorre, R. J. (2005). Attention to Simultaneous Unrelated Auditory and Visual Events: Behavioral and Neural Correlates. *Cerebral Cortex*, *15*, 1609–1620.

Jonides, J. (1981). Voluntary versus automatic control over the mind's eye's movement. In A. D. Baddeley & J. B. Long (Eds.), *Attention and performance IX*, (Vol. 9, pp. 187-203). Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1999). Increased Activity in Human Visual Cortex during Directed Attention in the Absence of Visual Stimulation. *Neuron*, *22*, 751–761.

Kastner, S., & Ungerleider, L. G. (2000). Mechanisms of visual attention in the human cortex. *Annual Review of Neuroscience*, *23*, 315–341.

Katus, T., Andersen, S. K., & Müller, M. M. (2012a). Maintenance of tactile short-term memory for locations is mediated by spatial attention. *Biological Psychology, 89*, 39–46.

Katus, T., Andersen, S. K., & Müller, M. M. (2012b). Nonspatial cueing of tactile STM causes shift of spatial attention. *Journal of Cognitive Neuroscience, 24*, 1596–1609.

Katus, T., Andersen, S. K., & Müller, M. M.(2014). Common mechanisms of spatial attention in memory and perception: A tactile dual-task study. *Cerebral Cortex, 24*, 707–718.

Kelly, S. P., Lalor, E. C., Reilly, R. B., & Foxe, J. J. (2006). Increases in Alpha Oscillatory Power Reflect an Active Retinotopic Mechanism for Distracter Suppression During Sustained Visuospatial Attention. *Journal of Neurophysiology*, *95*, 3844–3851.

Kiss, M., Grubert, A., Petersen, A., & Eimer, M. (2012). Attentional Capture by Salient Distractors during Visual Search Is Determined by Temporal Task Demands. *Journal of Cognitive Neuroscience*, *24*(3), 749–759.

Klimesch, W. (2011). Evoked alpha and early access to the knowledge system: The P1 inhibition timing hypothesis. *Brain Research*, *1408*, 52–71.

Klimesch, W., Vogt, F., & Doppelmayr, M. (1999). Interindividual differences in alpha and theta power reflect memory performance. *Intelligence, 27(4)*, 347-362.

Klinzing, J. G., Niethard, N., & Born, J. (2019). Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience*, *22*, 1596–1610.

Konkel, A., & Cohen, N. J. (2009). Relational memory and the hippocampus: representations and methods. *Frontiers in Neuroscience*, *3*(2), 166–174.

Konkoly, K.R., Appel, K., Chabani, E., Mironov, A. Y., Mangiaruga, A., Gott, J., Mallett, R., Caughran, B., Witkowski, S., Whitmore, N., Berent, J., Weber, F., Pipa, G., Türker, B., Maranci, J.-B., Sinin, A., Dorokhov, V., Arnulf, I., Oudiette, D., Dresler, M., & Paller, K. (2020). Real-time dialogue between experimenters and dreamers during REM sleep. *Work. Pap.,* Northwestern Univ., Evanston, IL.

Kolb, B., and Wishaw, I.Q. (1996). *Fundamentals of Human Neuropsychology.* New York: W.H. Freeman and Company.

Koster, E. H. W., Fox, E., & MacLeod, C. (2009). Introduction to the Special Section on Cognitive Bias Modification in Emotional Disorders. *Journal of Abnormal Psychology*, *118*(1), 1–4.

Kuhl, B. A., Bainbridge, W. A., & Chun, M. M. (2012). Neural Reactivation Reveals Mechanisms for Updating Memory. *The Journal of Neuroscience*, *32*(10), 3453–3461.

Kuhl, B. A., & Chun, M. M. (2014). Memory and Attention. In A. C. Nobre, & S. Kastner (Eds.), *The Oxford Handbook of Attention* (pp. 806-836). Oxford: Oxford University Press.

Kuhl, B. A., Rissman, J., Chun, M. M., & Wagner, A. D. (2011). Fidelity of neural reactivation reveals competition between memories. *PNAS*, *108*(14), 5903–5908.

Kuo, B.-C., Stokes, M. G., & Nobre, A. C. (2012). Attention Modulates Maintenance of Representations in Visual Short-term Memory. *Journal of Cognitive Neuroscience*, *24*(1), 51–60.

Kuo, B.-C., Stokes, M. G., Murray, A. M., & Nobre, A. C. (2014). Attention Biases Visual Activity in Visual Short-term Memory. *Journal of Cognitive Neuroscience*, *26*(7), 1377–1389.

Landman, R., Spekreijse, H., & Lamme, V. A. F. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, *43*, 149–164.

LaRocque, J. J., Eichenbaum, A. S., Starrett, M. J., Rose, N. S., Emrich, S. M., & Postle, B. R. (2015). The short- and long-term fates of memory items retained outside the focus of attention. *Memory & Cognition*, *43*(3), 453–468.

LaRocque, J. J., Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2013). Decoding attended information in short-term memory: An EEG study. *Journal of Cognitive Neuroscience*, *25*(1), 127–142.

LaRocque, J. J., Riggall, A. C., Emrich, S. M., & Postle, B. R. (2017). Within-Category Decoding of Information in Different Attentional States in Short-Term Memory. *Cerebral Cortex*, *27*, 4881–4890.

Latchoumane, C.-F. V, Ngo, H. V, Born, J., & Shin, H.-S. (2017). Thalamic Spindles Promote Memory Formation during Sleep through Triple Phase-Locking of Cortical, Thalamic, and Hippocampal Rhythms. *Neuron*, *95*(2), 424–435.

Leber, A. B., & Egeth, H. E. (2006). It's under control: Top-down search strategies can override attentional capture. *Psychological Bulletin & Review*, *13*(1), 132–138.

Leclercq, Y., Schrouff, J., Noirhomme, Q., Maquet, P., & Phillips, C. (2011). fMRI Artefact rejection and Sleep Scoring Toolbox. *Computational Intelligence and Neuroscience,* ID: 598206, 1-11.

Lepsien, J., & Nobre, A. C. (2007). Attentional Modulation of Object Representations in Working Memory. *Cerebral Cortex*, *17*, 2072–2083.

Lepsien, J., Thornton, I., & Nobre, A. C. (2011). Modulation of working-memory maintenance by directed attention. *Neuropsychologia*, *49*(6), 1569-1577.

Leszczyński, M., & Staudigl, T. (2016). Memory-guided attention in the anterior thalamus. *Neuroscience and Biobehavioral Reviews*, *66*, 163–165.

Lewis, P. A., & Durrant, S. J. (2011). Overlapping memory replay during sleep builds cognitive schemata. *Trends in Cognitive Sciences*, *15*(8), 343–351.

Lewis-Peacock, J. A., & Postle, B. R. (2008). Temporary activation of long-term memory supports working memory. *Journal of Neuroscience*, *28*(35), 8765–8771.

Liesefeld, H. R., & Müller, H. J. (2019). Distractor handling via dimension weighting. *Current Opinion in Psychology*, *29*, 160–167.

Liu, Y., Dolan, R. J., Kurth-Nelson, Z., & Behrens, T. E. J. (2019). Human Replay Spontaneously Reorganizes Experience. *Cell*, *178*(3), 640–652.

Loftus, E. F. (2005). Planting misinformation in the human mind: A 30-year investigation of the malleability of memory. *Learning & Memory*, *12*, 361–366.

Loftus, E. F., & Palmer, J. C. (1974). Reconstruction of Automobile Destruction: An Example of the Interaction Between Language and Memory. *Journal of Verbal Learning and Verbal Behavior*, *13*, 585–589.

Luck, Steven J. (2005). *An Introduction to the Event-Related Potential Technique.* Cambridge, Mass: MIT Press.

Luck, S. J., & Hillyard, S. A. (1994). Electrophysiological correlates of feature analysis during visual search. *Psychophysiology*, *31*, 291–308.

Luck, S. J., & Kappenman, E. S. (2012). ERP components and selective attention. In S. J. Luck & E. S. Kappenman (Eds.), *The Oxford Handbook of Event-Related Potential Components* (p. 295–327). Oxford University Press.

Makovski, T., & Pertzov, Y. (2015). Attention and memory protection: Interactions between retrospective attention cueing and interference. *The Quarterly Journal of Experimental Psychology*, *68*(9), 1735–1743.

Malhotra, R. K. & Avidan, A. Y. (2013). Sleep stages and scoring technique. In S. Chokrovetry & R. J. Thomas (eds.). *Atlas of Sleep Medicine*. Philadelphia, PA.

Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I . Role of features. *Memory & Cognition*, *22*(6), 657–672.

Maquet, P. (2001). The Role of Sleep in Learning and Memory. *Science*, *294*, 1048–1052.

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG-and MEG-data. *Journal of neuroscience methods*, *164*(1), 177-190.

Markant, J., & Amso, D. (2013). Selective memories: Infants' encoding is enhanced in selection via suppression. *Developmental Science*, *16*(6), 926–940.

Markant, J., Worden, M. S., & Amso, D. (2015). Not all attention orienting is created equal: Recognition memory is enhanced when attention orienting involves distractor suppression. *Neurobiology of Learning and Memory*, *120*, 28–40.

Marr, D. (1971). Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society of London. Biological Sciences, 262,* 23-81.

Marshall, L., & Born, J. (2007). The contribution of sleep to hippocampus-dependent memory consolidation. *Trends in Cognitive Sciences*, *11*(10), 442–450.

Martín-Buro, M. C., Wimber, M., Henson, R. N., & Staresina, B. P. (2020). Alpha Rhythms Reveal When and Where Item and Associative Memories Are Retrieved. *The Journal of Neuroscience*, *40*(12), 2510–2518.

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

McDonald, J. J., Green, J. J., Jannati, A., & Di Lollo, V. (2013). On the Electrophysiological Evidence for the Capture of Visual Attention. *Journal of Experimental Psychology. Human Perception and Performance*, *39*(3), 849–860.

Meister, M. L. R., & Buffalo, E. A. (2016). Getting directions from the hippocampus: The neural connection between looking and memory. *Neurobiology of Learning and Memory*, *134*, 135–144.

Milner, B., Corkin, S., & Teuber, H.-L. (1968). Further analysis of the hippocampal amnesic syndrome: 14-year follow-up study of H.M. *Neuropsychologia*, *6*, 215–234.

Moher, J., & Egeth, H. E. (2012). The ignoring paradox: Cueing distractor features leads first to selection, then to inhibition of to-be-ignored items. *Attention, Perception, and Psychophysics*, *74*(8), 1590–1605.

Moray, N. (1959). Attention in dichotic listening: Affective cues and the influence of instructions. *The Quarterly Journal of Experimental Psychology, 11*, 56-60.

Moscovitch, M. (2008). The Hippocampus As a "Stupid," Domain-Specific Module: Implications for Theories of Recent and Remote Memory, and of Imagination. *Canadian Journal of Experimental Psychology*, *62*(1), 62–79.

Moscovitch, M., Cabeza, R., Winocur, G., & Nadel, L. (2016). Episodic Memory and Beyond: The Hippocampus and Neocortex in Transformation. *Annual Review of Psychology*, *67*, 105–134.

Müller, H. J., & Krummenacher, J. (2006). Visual search and selective attention. *Visual Cognition*, *14*(4–8), 389–410.

Myers, N. E., Chekroud, S. R., Stokes, M. G., & Nobre, A. C. (2018). Benefits of Flexible Prioritization in Working Memory Can Arise Without Costs. *Journal of Experimental Psychology*, *44*(3), 398–411.

Myers, N. E., Stokes, M. G., & Nobre, A. C. (2017). Prioritizing Information during Working Memory: Beyond Sustained Internal Attention. *Trends in Cognitive Sciences*, *21*(6), 449–461.

Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., & Buzsáki, G. (1999). Replay and Time Compression of Recurring Spike Sequences in the Hippocampus. *The Journal of Neuroscience*, *19*(21), 9497–9507.

Nadel, L., & Moscovitch, M. (1997). Memory consolidation, retrograde amnesia and the hippocampal complex. *Current Opinion in Neurobiology, 7(2),* 217-227.

National Collaborating Centre for Mental Health (UK). (2005). *Post-Traumatic Stress Disorder: The Management of PTSD in Adults and Children in Primary and Secondary Care*. Gaskell.

Nee, D. E., & Jonides, J. (2009). Common and Distinct Neural Correlates of Perceptual and Memorial Selection. *NeuroImage*, *45*(3), 963–975.

Nelissen, N., Stokes, M., Nobre, A. C., & Rushworth, M. F. S. (2013). Frontal and Parietal Cortical Interactions with Distributed Visual Representations during Selective Attention and Action Selection. *The Journal of Neuroscience*, *33*(42), 16443–16458.

Neely, J.H., 1991. Semantic priming effects in visual word recognition: a selective review of current findings and theories. In E. Besner, & G. W. Humphreys (Eds.), *Basic Processes in Reading: Visual Word Recognition* (pp. 264–336). Hillsdale, NJ: Erlbaum.

Neisser, U. (1967). *Cognitive psychology*. New York: Appleton-Century-Crofts.

Neumann, O. (1984). Automatic processing: A review of recent findings and a plea for an old theory. In W. Prinz & A. F. Sanders (Eds.), *Cognition and Motor Processes* (pp. 255–290). Berlin: Springer-Verlag.

Neylan, T. C., Lenoci, M., Rothlind, J., Metzler, T. J., Schuff, N., Du, A.-T., Franklin, K. W., Weiss, D. S., Weiner, M. W., & Marmar, C. R. (2004). Attention, Learning, and Memory in Posttraumatic Stress Disorder. *Journal of Traumatic Stress*, *17*(1), 41–46.

Nobre, A. C. (2018). Attention. In J. T. Wixted, & J. Serences (eds). *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience, 4th Edition*. Hoboken, New Jersey: John Wiley & Sons, Inc.

Nobre, A. C., Coull, J. T., Maquet, P., Frith, C. D., Vandenberghe, R., & Mesulam, M. M. (2004). Orienting attention to locations in perceptual versus mental representations. *Journal of Cognitive Neuroscience*, *16*(3), 363–373.

Nobre, A. C., & Mesulam, M.-M. (2014). Large-scale networks for attentional biases. In A. C. Nobre & S. Kastner (eds.). *The Oxford Handbook of Attention* (pp. 105-151). Oxford: Oxford University Press.

Nobre, A. C., & Stokes, M. G. (2019). Premembering Experience: A Hierarchy of Time-Scales for Proactive Attention. *Neuron*, *104*(1), 132–146.

Nobre, A. C., & Stokes, M. G. (2020). Memory and attention: the back and forth. In M. S. Gazzaniga (ed). *The Cognitive Neurosciences*. Cambridge: MIT Press.

Noonan, M., Adamian, N., Pike, A., Printzlau, F., Crittenden, B. M., & Stokes, M. G. (2016). Distinct mechanisms for distractor suppression and target facilitation. *The Journal of Neuroscience*, *36*(6), 1797–1807.

Nordfang, M., & Wolfe, J. M. (2018). Guided search through memory. *Visual Cognition*, *26*(4), 285–298.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford: Clarendon Press.

Olivers, C. N. L., Meijer, F., & Theeuwes, J. (2006). Feature-Based Memory-Driven Attentional Capture: Visual Working Memory Content Affects Visual Attention. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(5), 1243–1265.

Ong, J. L., Lo, J. C., Chee, N. I. Y. N., Santostasi, G., Paller, K. A., Zee, P. C., & Chee, M. W. L. (2016). Effects of phase-locked acoustic stimulation during a nap on EEG spectra and declarative memory consolidation. *Sleep Medicine*, *20*, 88–97.

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J.-M. (2011). FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. *Computational Intelligence and Neuroscience*, *2011*, 156869.

Oudiette, D., Antony, J. W., Creery, J. D., & Paller, K. A. (2013). The Role of Memory Reactivation during Wakefulness and Sleep in Determining Which Memories Endure. *Journal of Neuroscience*, *33*(15), 6672–6678.

Oudiette, D., & Paller, K. A. (2013). Upgrading the sleeping brain with targeted memory reactivation. *Trends in Cognitive Sciences*, *17*(3), 142–149.

Paller, K. A., Creery, J. D., & Schechtman, E. (2021). Memory and Sleep: How Sleep Cognition Can Change the Waking Mind for the Better. *Annual Review of Psychology*, *72*(18), 1–28.

Park, H., Lee, D. S., Kang, E., Kang, H., Hahm, J., Kim, J. S., Chung, C. K., & Jensen, O. (2014). Blocking of irrelevant memories by posterior alpha activity boosts memory encoding. *Human Brain Mapping*, *35*(8), 3972–3987.

Patai, E. Z., Doallo, S., & Nobre, A. C. (2012). Long-term Memories Bias Sensitivity and Target Selection in Complex Scenes. *Journal of Cognitive Neuroscience*, *24*(12), 2281–2291.

Patel, S. H., & Azzam, P. N. (2005). Characterization of N200 and P300: Selected studies of the Event-Related Potential. *International Journal of Medical Sciences*, *2*(4), 147–154.

Perrin, F., Pernier, J., Bertrand, O., & Echallier, J. F. (1989). Spherical splines for scalp potential and current density mapping. *Electroencephalography and clinical neurophysiology*, *72*(2), 184-187.

Pertzov, Y., Bays, P. M., Joseph, S., & Husain, M. (2013). Rapid forgetting prevented by retrospective attention cues. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(5), 1224–1231.

Pertzov, Y., Manohar, S., & Husain, M. (2017). Rapid Forgetting Results From Competition Over Time Between Items in Visual Working Memory. *Journal of Experimental Psychology*, *43*(4), 528–536.

Pilon, M., Desautels, A., Montplaisir, J., & Zadra, A. (2012). Auditory arousal responses and thresholds during REM and NREM sleep of sleepwalkers and controls. *Sleep Medicine*, *13*(5), 490–495.

Poch, C., Campo, P., & Barnes, G. R. (2014). Modulation of alpha and gamma oscillations related to retrospectively orienting attention within working memory. *European Journal of Neuroscience*, *40*(2), 2399–2405.

Posner, M. I. (1978). *Chronometric explorations of mind*. Oxford: Erlbaum.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*(1), 3–25.

Rasch, B., & Born, J. (2013). About Sleep's Role in Memory. *Physiological Review*, *93*, 681–766.

Rasch, B., Christian, B., Gais, S., & Born, J. (2007). Odor Cues During Slow-Wave Sleep Prompt Declarative Memory Consolidation. *Science*, *315*(5817), 1426–1429.

Richter, F. R., Cooper, R. A., Bays, P. M., & Simons, J. S. (2016). Distinct neural mechanisms underlie the success, precision, and vividness of episodic memory. *eLife*, *5*, 1–18.

Rolls, E. T. (2013). The mechanisms for pattern completion and pattern separation in the hippocampus. *Frontiers in Systems Neuroscience*, *7*, 1–21.

Rosen, M. L., Stern, C. E., Devaney, K. J., & Somers, D. C. (2018). Cortical and Subcortical Contributions to Long-Term Memory-Guided Visuospatial Attention. *Cerebral Cortex*, *28*, 2935–2947.

Rosen, M. L., Stern, C. E., Michalka, S. W., Devaney, K. J., & Somers, D. C. (2015a). Influences of Long-Term Memory-Guided Attention and Stimulus-Guided Attention on Visuospatial Representations within Human Intraparietal Sulcus. *The Journal of Neuroscience*, *35*(32), 11358–11363.

Rosen, M. L., Stern, C. E., Michalka, S. W., Devaney, K. J., & Somers, D. C. (2015b). Cognitive Control Network Contributions to Memory-Guided Visual Attention. *Cerebral Cortex*, 1–15.

Rosen, M. L., Stern, C. E., & Somers, D. C. (2014). Long-term memory guidance of visuospatial attention in a change-detection paradigm. *Frontiers in Psychology*, *5*(266), 1–8.

Rudoy, J. D., Voss, J. L., Westerberg, C. E., & Paller, K. A. (2009). Strengthening Individual Memories by Reactivating Them During Sleep. *Science*, *326*(November), 1079.

Rugg, M. D. (1995). Event-related potential studies of human memory. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (p. 789–801). The MIT Press.

Salvato, G., Patai, E. Z., McCloud, T., & Nobre, A. C. (2016a). Apolipoprotein 4 breaks the association between declarative long-term memory and memory-based orienting of spatial attention in middle-aged individuals. *Cortex*, *82*, 206–216.

Salvato, G., Patai, E. Z., & Nobre, A. C. (2016b). Preserved memory-based orienting of attention with impaired explicit memory in healthy ageing. *Cortex*, *74*, 67–78.

Santostasi, G., Malkani, R., Riedner, B., Bellesi, M., Tononi, G., Paller, K. A., & Zee, P. C. (2016). Phase-locked loop for precisely timed acoustic stimulation during sleep. *Journal of Neuroscience Methods*, *259*, 101–114.

Sauter, M., Liesefeld, H. R., Zehetleitner, M., & Müller, H. J. (2018). Region-based shielding of visual search from salient distractors: Target detection is impaired

with same- but not different-dimension distractors. *Attention, Perception, & Psychophysics*, *80*, 622–642.

Sawaki, R., & Luck, S. J. (2010). Capture versus suppression of attention by salient singletons: Electrophysiological evidence for an automatic attend-to-me signal. *Attention, Perception & Psychophysics*, *72*(6), 1455–1470.

Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *The Journal of Neuropsychiatry and Clinical Neurosciences*, *20*(11), 11–21.

Schechtman, E., Witkowski, S., Lampe, A., Wilson, B. J., & Paller, K. A. (2020). Targeted memory reactivation during sleep boosts intentional forgetting of spatial locations. *Scientific Reports*, *10*(2327), 1–9.

Schönauer, M., Grätsch, M., & Gais, S. (2015). Evidence for two distinct sleep-related long-term memory consolidation processes. *Cortex*, *63*, 68–78.

Schouten, D. I., Pereira, 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.

Schrouff, J., Coppieters, D., Lehembre, R., Leclercq, Y., Maquet, P., Noirhomme, Q., Degueldre, C., Balteau, E., & Phillips, C. (2009). fMRI Artefact rejection and Sleep Scoring Toolbox (FASST). Cyclotron research Centre, University of Liège, Belgium.

Skaggs, W. E., & McNaughton, B. L. (1996). Replay of Neuronal Firing Sequences in Rat Hippocampus During Sleep Following Spatial Experience. *Science*, *271*, 1870–1873.

Slagter, H. A., Prinssen, S., Reteig, L. C., & Mazaheri, A. (2016). Facilitation and inhibition in attention: Functional dissociation of pre-stimulus alpha activity, P1, and N1 components. *NeuroImage*, *125*, 25–35.

Sligte, I. G., Scholte, H. S., & Lamme, V. A. (2009). V4 activity predicts the strength of visual short-term memory representations. *Journal of Neuroscience*, *29*(23), 7432-7438.

Soto, D., Heinke, D., Humphreys, G. W., & Blanco, M. J. (2005). Early, Involuntary Top-Down Guidance of Attention From Working Memory. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(2), 248–261.

Souza, A. S., & Oberauer, K. (2016). In search of the focus of attention in working memory: 13 years of the retro-cue effect. *Attention, Perception, & Psychophysics*, *78*, 1839–1860.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, *74*, 1–29.

Squire, L. R., & Knowlton, B. J. (1995). Memory, hippocampus, and brain systems. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (p. 825–837). The MIT Press.

Squire, L. R., Stark, C. E. L., & Clark, R. E. (2004). The Medial Temporal Lobe. *Annual Review of Neuroscience*, *27*, 279–306.

Stallone, A., Cicone, A., & Materassi, M. (2020). New insights and best practices for the successful use of Empirical Mode Decomposition, Iterative Filtering and derived algorithms. *Scientific Reports*, *10*(15161), 1–15.

Staresina, B. P., & Wimber, M. (2019). A Neural Chronometry of Memory Recall. *Trends in Cognitive Sciences*, *23*(12), 1071–1085.

Steriade, M., McCormick, D. A., & Sejnowski, T. J. (1993). Thalamocortical Oscillations in the Sleeping and Aroused Brain. *Science*, *262*, 679–685.

Stickgold, R. (2005). Sleep-dependent memory consolidation. *Nature*, *437*, 1272–1278.

Stickgold, R. (2013). Parsing the role of sleep in memory processing. *Current Opinion in Neurobiology*, *23*, 847–853.

Stokes, M. G., Atherton, K., Patai, E. Z., & Nobre, A. C. (2012). Long-term memory prepares neural activity for perception. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(6), E360-7.

Stokes, M., Thompson, R., Nobre, A. C., & Duncan, J. (2009). Shape-specific preparatory activity mediates attention to targets in human visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(46), 19569–19574.

Stropahl, M., Bauer, A. K. R., Debener, S., & Bleichner, M. G. (2018). Source-Modeling auditory processes of EEG data using EEGLAB and brainstorm. *Frontiers in Neuroscience*, *12*(MAY), 1–11.

Strunk, J., Morgan, L., Reaves, S., Verhaeghen, P., & Duarte, A. (2019). Retrospective Attention in Short-Term Memory Has a Lasting Effect on Long-Term Memory Across Age. *Journals of Gerontolgy: Series B*, *74*(8), 1317-1325.

Summerfield, J. J., Lepsien, J., Gitelman, D. R., Mesulam, M. M., & Nobre, A. C. (2006). Orienting attention based on long-term memory experience. *Neuron*, *49*, 905–916.

Summerfield, J. J., Rao, A., Garside, N., & Nobre, A. C. (2011). Biasing Perception by Spatial Long-Term Memory. *Journal of Neuroscience*, *31*(42), 14952–14960.

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.

Tarder-Stoll, H., Jayakumar, M., Dimsdale-Zucker, H. R., Günseli, E., & Aly, M. (2020). Dynamic internal states shape memory retrieval. *Neuropsychologia*, *138*, 107328.

Theeuwes, J. (1990). Perceptual selectivity is task dependent: Evidence from selective search. *Acta psychologica*, *74*(1), 81-99.

Theeuwes, J. (2004). Top-down search strategies cannot override attentional capture. *Psychological Bulletin & Review*, *11*(1), 65–70.

Theeuwes, J. (2010). Top – down and bottom – up control of visual selection. *Acta Psychologica*, *135*(2), 77–99.

Theeuwes, J. (2018). Visual Selection: Usually Fast and Automatic; Seldom Slow and Volitional. *Journal of Cognition*, *1(1)*(29), 1–15.

Theeuwes, J. (2019). Goal-driven, stimulus-driven, and history-driven selection. *Current Opinion in Psychology*, *29*, 97–101.

Theeuwes, J., Kramer, A. F., Hahn, S., Irwin, D. E., & Zelinsky, G. J. (1999). Influence of attentional capture on oculomotor control. *Journal of Experimental Psychology: Human Perception and Performance*, *25*(6), 1595–1608.

Theeuwes, J., & Van der Burg, E. (2011). On the limits of top-down control of visual selection. *Attention, Perception, & Psychophysics*, *73*, 2092–2103.

Tibboel, H., & Liefooghe, B. (2020). Attention for future reward. *Psychological Research*, *84*, 706–712.

Todd, R. M., & Manaligod, M. G. M. (2018). Implicit guidance of attention: The priority state space framework. *Cortex*, *102*, 121–138.

Treisman, A. M. (1969). Strategies and models of selective attention. *Psychological review*, *76*(3), 282-299.

Tulving, E. (1983). Elements of episodic memory. Oxford University Press: New York.

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

Turatto, M., Bonetti, F., & Pascucci, D. (2018). Filtering visual onsets via habituation: A context-specific long-term memory of irrelevant stimuli. *Psychonomic Bulletin & Review*, *25*, 1028–1034.

Turatto, M., & Pascucci, D. (2016). Short-term and long-term plasticity in the visual-attention system: Evidence from habituation of attentional capture. *Neurobiology of Learning and Memory*, *130*, 159–169.

Turk-Browne, N. B., Yi, D. J., & Chun, M. M. (2006). Linking implicit and explicit memory: common encoding factors and shared representations. *Neuron*, *49*(6), 917-927.

Uddo, M., Vasterling, J. J., Brailey, K., & Sutker, P. B. (1993). Memory and Attention in Combat-Related Post-Traumatic Stress Disorder (PTSD). *Journal of Psychopathology and Behavioral Assessment*, *15*(1), 43–52.

Uncapher, M. R., Hutchinson, J. B., & Wagner, A. D. (2011). Dissociable Effects of Top-Down and Bottom-Up Attention during Episodic Encoding. *The Journal of Neuroscience*, *31*(35), 12613–12628.

Uncapher, M. R., & Rugg, M. D. (2009). Selecting for Memory? The Influence of Selective Attention on the Mnemonic Binding of Contextual Information. *The Journal of Neuroscience*, *29*(25), 8270–8279.

Uncapher, M. R., & Wagner, A. D. (2009). Posterior parietal cortex and episodic encoding: Insights from fMRI subsequent memory effects and dual-attention theory. *Neurobiology of Learning and Memory*, *91*(2), 139–154.

van der Helm, E., Gujar, N., Nishida, M., & Walker, M. P. (2011). Sleep-Dependent Facilitation of Episodic Memory Details. *PLoS ONE*, *6*(11), e27421:1-10.

van der Stigchel, S., & Theeuwes, J. (2006). Our eyes deviate away from a location where a distractor is expected to appear. *Experimental Brain Research*, *169*(3), 338–349.

van Ede, F., Chekroud, S. R., & Nobre, A. C. (2019). Human gaze tracks attentional focusing in memorized visual space. *Nature Human Behaviour*, *3*(5), 462–470.

van Moorselaar, D., Theeuwes, J., & Olivers, C. N. L. (2019). Memory-based attentional biases survive spatial suppression driven by selection history. *Visual Cognition*, *27*(3–4), 343–350.

Vargas, I. M., Schechtman, E., & Paller, K. A. (2019). Targeted memory reactivation during sleep to strengthen memory for arbitrary pairings. *Neuropsychologia*, *124*(December 2018), 144–150.

Vatterott, D. B., & Vecera, S. P. (2012). Experience-dependent attentional tuning of distractor rejection. *Psychological Bulletin & Review*, *18*, 871–878.

Wais, P. E., Rubens, M. T., Boccanfuso, J., & Gazzaley, A. (2010). Neural Mechanisms Underlying the Impact of Visual Distraction on Retrieval of Long-Term Memory. *The Journal of Neuroscience*, *30*(25), 8541–8550.

Waldhauser, G. T., Braun, V., & Hanslmayr, S. (2016). Episodic Memory Retrieval Functionally Relies on Very Rapid Reactivation of Sensory Information. *The Journal of Neuroscience*, *36*(1), 251–260.

Wallis, G., Stokes, M., Cousijn, H., Woolrich, M., & Nobre, A. C. (2015). Frontoparietal and Cingulo-opercular Networks Play Dissociable Roles in Control of Working Memory. *Journal of Cognitive Neuroscience*, *27*(10), 2019–2034.

Wamsley, E. J. (2019). Memory Consolidation during Waking Rest. *Trends in Cognitive Sciences*, *23*(3), 171–173.

Wang, B., Antony, J. W., Lurie, S., Brooks, P. P., Paller, K. A., & Norman, K. A. (2019). Targeted Memory Reactivation during Sleep Elicits Neural Signals Related to Learning Content. *The Journal of Neuroscience*, *39*(34), 6728–6736.

Wang, B., Samara, I., & Theeuwes, J. (2019). Statistical regularities bias overt attention. *Attention, Perception, and Psychophysics*, *81*(6), 1813–1821.

Wang, B., & Theeuwes, J. (2018a). Statistical regularities modulate attentional capture independent of search strategy. *Attention, Perception, & Psychophysics*, *80*, 1763–1774.

Wang, B., & Theeuwes, J. (2018b). How to inhibit a distractor location? Statistical learning versus active, top-down suppression. *Attention, Perception, and Psychophysics*, *80*(4), 860–870.

Wascher, E., & Wolber, M. (2004). Attentional and intentional cueing in a Simon task: An EEG-based approach. *Psychological Research*, *68*, 18–30.

Westerberg, C. E., Florczak, S. M., Weintraub, S., Mesulam, M., Marshall, L., Zee, P. C., & Paller, K. A. (2015). Memory improvement via slow-oscillatory stimulation during sleep in older adults. *Neurobiology of Aging*, *36*(9), 2577–2586.

Westermann, J., Lange, T., Textor, J., & Born, J. (2015). System Consolidation During Sleep – A Common Principle Underlying Psychological and Immunological Memory Formation. *Trends in Neurosciences*, *38*(10), 585–597.

Wilhelm, I., Wagner, U., & Born, J. (2011a). Opposite Effects of Cortisol on Consolidation of Temporal Sequence Memory during Waking and Sleep. *Journal of Cognitive Neuroscience*, *23*(12), 3703–3712.

Wilhelm, I., Diekelmann, S., Molzow, I., Ayoub, A., Mölle, M., & Born, J. (2011b). Sleep Selectively Enhances Memory Expected to Be of Future Relevance. *The Journal of Neuroscience*, *31*(5), 1563–1569.

Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of Hippocampal Ensemble Memories During Sleep. *Science*, *265*(July), 676–679.

Wimber, M., Alink, A., Charest, I., Kriegeskorte, N., & Anderson, M. C. (2015). Retrieval induces adaptive forgetting of competing memories via cortical pattern suppression. *Nature Neuroscience*, *18*(4), 582–589.

Wimber, M., Maaß, A., Staudigl, T., Richardson-Klavehn, A., & Hanslmayr, S. (2012). Rapid memory reactivation revealed by oscillatory entrainment. *Current Biology*, *22*(16), 1482–1486.

Wimmer, R. D., Schmitt, L. I., Davidson, T. J., Nakajima, M., Deisseroth, K., & Halassa, M. M. (2015). Thalamic control of sensory selection in divided attention. *Nature*, *526*(7575), 705–709.

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

Woud, M. L., Verwoerd, J., & Krans, J. (2017). Modification of cognitive biases related to posttraumatic stress: A systematic review and research agenda. *Clinical Psychology Review*, *54*, 81–95.

Wundt, W. M. (1874). *Grundzüge der physiologischen Psychologie* (Vol. 1). W. Engelman.

Wixted, J. T. (2018). *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience*. John Wiley & Sons.

Yantis, S., & Jonides, J. (1984). Abrupt Visual Onsets and Selective Attention: Evidence From Visual Search. *Journal of Experimental Psychology: Human Perception and Performance*, *10*(5), 601–621.

Yerkes, R. M., & Dodson, J. D. (1908). The relation of strength of a stimulus to rapidity of habit-formation. *Journal of Comparative Neurology and Psychology, 18*(5), 459–482.

Yi, D.-J., & Chun, M. M. (2005). Attentional Modulation of Learning-Related Repetition Attenuation Effects in Human Parahippocampal Cortex. *The Journal of Neuroscience*, *25*(14), 3593–3600.

Yu, X., & Geng, J. J. (2019). The attentional template is shifted and asymmetrically sharpened by distractor context. *Journal of Experimental Psychology: Human Perception and Performance, 45*(3), 336–353.

Zhang, H., Fell, J., & Axmacher, N. (2018). Electrophysiological mechanisms of human memory consolidation. *Nature Communications*, *9*(4103), 1–11.

Zokaei, N., Manohar, S., Husain, M., & Feredoes, E. (2014). Causal Evidence for a Privileged Working Memory State in Early Visual Cortex. *The Journal of Neuroscience*, *34*(1), 158–162.