How event-based memories change as a function of forgetting and consolidation

Bardur Hofgaard Joensen

Doctor of Philosophy

University of York

Psychology

December 2019

Abstract

Episodic memories are composed of multiple elements, from the people we encounter, the locations we visit, and the objects we interact with. These 'episodes' are thought to be stored in memory as coherent event representations and are associated with holistic recollection at retrieval, such that the retrieval of one element is dependent on the retrieval of all other elements from the same event. Evidence for this 'dependency' has been shown to emerge from the association between the event elements themselves. Critically, dependency is seen when participants learn three overlapping pairwise associations in a 'closed-loop', but not when participants learn only two out of the three possible associations in an 'open-loop', suggesting that all pairwise associations between event-elements need to be explicitly encoded for a coherent event representation to emerge.

Here I asked whether the associative structure formed at encoding affects how event-based memories are retained over a period of forgetting and consolidation. Recently formed representations are susceptible to forgetting via interference and/or decay, but also undergo memory consolidation; becoming less susceptible to interference and/or decay. As such, retention for an event-based representation will reflect an interaction between forgetting and consolidation. This thesis presents evidence that closed-loops tend to be forgotten in an all-or-none manner, such that closed-loops are more likely to either be retained or forgotten in their entirety (Chapter 2). In contrast, open-loops are associated with a more asymmetrical pattern of forgetting as a function of memory reactivation during sleep (Chapter 3). Further, the thesis presents fMRI evidence that closed-loops continue to be retrieved in a coherent manner following a period of forgetting and consolidation (Chapter 4). These findings suggest that the associative structure formed at encoding has a lasting impact on the coherence of the underlying memory representation.

List of Contents

List of Tables

Chapter 2

Chapter 3

Chapter 4

List of Figures

Chapter 2

Chapter 3

Chapter 4

Acknowledgements

First and foremost, I would like to thank my supervisors Aidan Horner and Gareth Gaskell for the support and guidance they have provided over the past years. I wish to give a special thanks to my primary supervisor, Aidan Horner, for his encouragement to pursue new ideas and opportunities, for always making time to discuss my work with me, and for helping me develop as a scientist. I have been fortunate to have you as my primary supervisor and have truly enjoyed my time working with you.

Thanks also go to the members of the York Episodic Memory Lab (Emma James, Jamie Cockcroft, and Sam Berens), especially Sam for his tolerance whilst sharing an office with me for the past 3 years.

Lastly, I would also like to thank my family for their unconditional support, and my partner, Saana, whose kindness and encouragement has been invaluable throughout the past years. For that and more, I am truly grateful.

Author’s declaration

This thesis is the presentation of original work completed solely by the author, under the supervision of Dr Aidan Horner and Prof Gareth Gaskell. This work has not previously been presented for an award at this or any other university. All sources are acknowledged as References.

Publications:

The experiments presented in Chapter 2 have previously been reported in the following journal article.

Joensen, B. H., Gaskell, M. G., & Horner. A. J. (2019). United we fall: All-or-None Forgetting of Complex Episodic Events. *Journal of Experimental Psychology: General*. *149*(2), 230-248

1. General introduction

The concept of episodic memory refers to a distinct form of declarative memory for personally experienced events (Tulving, 1983). Episodic memory is proposed to be autonoetic and recollective in nature, allowing for the complete re-experience of previous life events (Tulving, 1983). Tulving (1983) proposed that episodic memory is the core function that allows us to re-immerse ourselves in our own personal past, providing the basis for our sense of self (Conway, 2005). This ability for 'mental time travel' (Tulving, 1983) is also thought to underlie our ability to make predictions about future events (Schacter, Addis, & Buckner, 2007).

Amnesia - the failure to acquire and retain episodic information (Tulving, 1983) – has been observed following damage to different regions of the brain, such as the frontal lobes (Tanji, Suzuki, Fujii, Higano, & Yamadori, 2003). However, amnesia is most commonly associated with damage to the medial temporal lobes (Aggleton & Brown, 1999; Cohen & Eichenbaum, 1993). Evidence for this has come from neuropsychological studies demonstrating that patients with bilateral damage to medial temporal lobe regions, such as the hippocampus, present with amnesia (Scoville & Milner, 1957; Vargha-Khadem et al., 1997; Woods, Schoene, & Kneisley, 1982). Notably, patients with selective damage to the hippocampus show marked deficits on tasks that assess performance for associative or relational memory (Hannula, Tranel, & Cohen, 2006; Holdstock, Mayes, Gong, Roberts, & Kapur, 2005; Kumaran et al., 2007; Mayes et al., 2004; Ryan, Althoff, Whitlow, & Cohen, 2000), consistent with the proposal that the role of the hippocampus in episodic memory is reflected in its capacity to bind different event elements (e.g., people, locations, and objects) into an 'event engram' (Damasio, 1989; Marr, 1971; McClelland, McNaughton, & O'Reilly, 1995; Teyler & DiScenna, 1986).

1.1. Medial temporal lobes and episodic memory

1.1.1. Representations in the medial temporal lobes

Regions within the medial temporal lobe, such as the perirhinal cortex and hippocampus, have been implicated in episodic memory (Squire & Zola-Morgan, 1991). However, whereas the perirhinal cortex is thought to support the representation of item information (Aggleton & Brown, 1999; Diana, Yonelinas, & Ranganath, 2007), the hippocampus has been shown to act as a convergence zone (Backus, Bosch, Ekman, Grabovetsky, & Doeller, 2016) for representing the association between different sources of information (Cohen et al., 1999; Davachi, 2006; Mayes, Montaldi, & Migo, 2007; Ranganath, 2010), such as the association between an object and the context in which it was studied (Davachi, Mitchell, & Wagner, 2003; Ranganath et al., 2004; Staresina & Davachi, 2006; Uncapher, Otten, & Rugg, 2006).

Consistent with the proposal that the perirhinal cortex supports memory for item information, previous work has demonstrated that increased BOLD activation in the perirhinal cortex for novel items at encoding (Pihlajamäki et al., 2003, 2004) predicts later item recognition (Awipi & Davachi, 2008; Staresina & Davachi, 2008). However, the perirhinal cortex has also been shown to support memory for perceptual and conceptual item features (O'Kane, Insler, & Wagner, 2005; Staresina & Davachi, 2006; Staresina, Duncan, & Davachi, 2011; Taylor, Moss, Stamatakis, & Tyler, 2006) and to aid the discrimination between items that share common features (Barense, Gaffan, & Graham, 2007). This has led to the proposal that the perirhinal cortex may not only play an important role in supporting memory for items, but also in representing the association between an item and its features (Bussey & Saksida, 2007).

Although the perirhinal cortex has been implicated in certain associative processes (Mayes et al., 2007; Watson, Wilding, & Graham, 2012), the hippocampus is thought to be distinct in its capacity to bind information across multiple domains (Cohen et al., 1999; Diana et al., 2007; Mayes et al., 2007). This is supported by evidence showing that encoding activity in the hippocampus predicts later memory success for objects and contexts (Cansino, Maquet, Dolan, & Rugg, 2002; Ranganath et al., 2004), in addition to studies showing that hippocampal activity at retrieval correlates with memory for associations between pairs of items and items and their associated context (Cansino et al., 2002; Fenker, Schott, Richardson-Klavehn, Heinze, & Düzel, 2005; Kirwan & Stark, 2004; Preston, Shrager, Dudukovic, & Gabrieli, 2004; Weis et al., 2004).

1.1.2. Retrieval processes in the hippocampus and perirhinal cortex

A distinction has also been made between the hippocampus and perirhinal cortex and whether they support memory via recollection or familiarity (Aggleton & Brown, 1999; Diana et al., 2007; Eichenbaum, Yonelinas, & Ranganath, 2007; Mayes et al., 2007; Yonelinas, 2002). Evidence has suggested that retrieval can be supported by two distinct processes – familiarity (e.g., the assessment of an item’s familiarity) and recollection (e.g., the retrieval of an item and the context in which it was encountered) - that are proposed to be functionally independent (see Yonelinas, 2002 for review; but see Wixted & Stretch, 2004).

Different retrieval dynamics for recollection and familiarity have been distinguished by applying retrieval operating characteristic (ROC) functions to memory performance (Yonelinas, 1994). ROC functions plot the correct identification of items previously studied against the incorrect identification of items as previously studied. Yonelinas (1994) showed that familiarity follows a symmetrical, curvilinear shape starting at zero, while the addition of recollection is characterised by an asymmetrical shape with an intercept greater than zero. This has been taken to suggest that while familiarity reflects a continuous memory signal denoting instances where an item is recognised or not, recollection is best characterised as a threshold-like retrieval process, where recollection either succeeds or fails (Yonelinas, 2002).

Differences in recollection and familiarity have commonly been assessed using a remember/know paradigm (Tulving, 1985). In this paradigm participants are required to give a judgement regarding their memory for an item; ‘remember’ if the participant, for instance, recalls both the item and its associated context; or ‘know’ if the participant only remembers encountering the item previously, but not the context in which it was learnt. As such, ‘remember’ judgements reflect the retrieval of both the item and any associative information, and ‘know’ judgements reflect memory for the item only. However, Donaldson (1996) argued that the distinction between remember/know judgements may simply reflect differences in memory strength, such that remember and know judgements are most accurately characterised as responses that fall above or below a given decision criteria.

Consistent with this, Ingram, Mickes, and Wixted (2012) showed that ‘remember’ and ‘know’ judgements can both be accompanied with high and low levels of confidence, suggesting that the strength of retrieval, as a function of recollection and/or familiarity, can differ in some continuous manner. This is inconsistent with the proposal that recollection reflects a threshold-like process where retrieval either succeeds or fails (Yonelinas, 1994), and provides support for the idea that both familiarity and recollection can be characterised by a continuous memory signal (Wixted & Stretch, 2004). However, Ingram et al. (2012) observed that even in instances where ‘remember’ judgements where made with low levels of confidence, the retrieval of a previously studied word was associated with higher levels of source memory accuracy (e.g., the colour or location of the word during encoding) than familiarity judgements made with high levels of confidence. This suggests that even when recollection is relatively ‘weak’, it still differs in some quantitative manner from familiarity. However, Kelley and Wixted (2001) argued that even when recollection succeeds, the contextual or associative information retrieved from memory will differ in some quantitative or qualitative way. This has led to the suggestion that recollection is most accurately characterised as a some-or-none process (Kelley & Wixted, 2001; Onyper, Zhang, & Howard, 2010). In this way, recollection may sometimes fail, and retrieval can only be supported by familiarity. However, even when recollection succeeds, the retrieved information can vary both in amount and quality (Harlow & Donaldson, 2013; Harlow & Yonelinas, 2016).

The roles of the hippocampus and perirhinal cortex in recollection and familiarity are captured in a computational account of recognition memory (Norman & O'Reilly, 2003; see also Marr, 1971; McClelland et al., 1995). Computational models have long proposed that the anatomical characteristics of the hippocampus make it particularly well-suited for the encoding and retrieval of event-based representations (Marr, 1971; McClelland et al., 1995; Norman & O'Reilly, 2003; Treves & Rolls, 1994). An important property of the hippocampus is that the representations within the dentate gyrus – a subregion of the hippocampus – are sparse (McClelland et al., 1995; Norman & O'Reilly, 2003). Sparse representations in the dentate gyrus are thought to support the process of 'pattern separation', allowing for the production of distinct and orthogonal representations of mnemonic information in the hippocampus (Bakker, Kirwan, Miller, & Stark, 2008; Berron et al., 2016; Kirwan & Stark, 2007; Yassa & Stark, 2011). At retrieval, recurrent connections in the cornu ammonis 3 (CA3) subregion of the hippocampus are thought to allow for the complete reconstruction of the original memory (Nakazawa et al., 2002; Neunuebel & Knierim, 2014) via a 'collateral effect' (Marr, 1971) or 'pattern completion' (Gardner-Medwin, 1976; Hopfield, 1982; McClelland et al., 1995; Treves & Rolls, 1992).

Note that this computational account (Norman & O'Reilly, 2003) assumes that the process of pattern completion (i.e., the complete retrieval of an entire pattern in the presence of a partial or ambiguous input) should lead to recollection that tends to be 'all-or-none', but it does not exclude the possibility that the information retrieved will differ in some quantitative or qualitative way. In contrast to the hippocampus, representations in the perirhinal cortex are less distinct (Norman & O'Reilly, 2003). The overlapping nature of representations in the perirhinal cortex is thought to constrain the retrieval of specific contextual details associated with an event. As such, the perirhinal cortex is only able to signal whether an item is familiar or not, based on the perceived item’s overlap with an existing representation (Norman & O’Reilly, 2003).

Consistent with this account, studies using the remember/know paradigm have made a distinction between recollection and familiarity in the perirhinal cortex and hippocampus, with increased hippocampal and perirhinal activity for items that are judged as remembered and known, respectively (Eldridge, Knowlton, Furmanski, Bookheimer, & Engel, 2000; Uncapher & Rugg, 2005; Wheeler & Buckner, 2004; Woodruff, Johnson, Uncapher, & Rugg, 2005). Further support for this distinction between the hippocampus and perirhinal cortex has come from neuroimaging studies showing that while activity in the hippocampus correlates with the retrieval of contextual information associated with a retrieved item, activity in the perirhinal cortex correlates with item recognition, but not the recollection of contextual information (Davachi et al., 2003; Ranganath et al., 2003). However, the proposal that the hippocampus is selectively involved in recollection (Aggleton & Brown, 1999; Yonelinas, 1994) has been challenged by evidence showing that patients with damage to the hippocampus appear to be equally impaired on tasks requiring recollection and familiarity judgements (Wais, Wixted, Hopkins, & Squire, 2006; Wixted & Stretch, 2004). This has motivated the proposal that the hippocampus may contribute to both recollection and familiarity, perhaps via a continuous rather than threshold-like process (Wais et al., 2006; but see Yonelinas, Aly, Wang, & Koen, 2010).

1.1.3. Summary

Despite uncertainty regarding how to best characterise the role of the hippocampus in recollection and familiarity (and indeed whether recollection and familiarity reflect distinct retrieval processes (see Wixted, 2007 for review)), evidence has accumulated to demonstrate that the hippocampus may support the encoding and retrieval of event-based representations in a manner that is distinct from other regions of the medial temporal lobes, such as the perirhinal cortex (Aggleton & Brown, 1999; Cohen & Eichenbaum, 1993; Davachi, 2006; Mayes et al., 2007). The distinct role of the hippocampus is captured in computational accounts demonstrating that the hippocampus is ideally suited for representing distinct events (via pattern separation) that can be reconstructed during later retrieval via the process of pattern completion (Marr, 1971; McClelland et al., 1995; Norman & O'Reilly, 2003; Treves & Rolls, 1994).

1.2. The hippocampus and holistic retrieval

A hallmark of episodic memory is that retrieval is holistic in nature (Tulving, 1983). For example, if we met a friend in a café and they gave us a present, reminiscing about this event should lead to the retrieval of all the elements (i.e., people, locations, and objects) of the event. Although this proposal is supported by computational accounts suggesting that recollection is all-or-none (Norman & O'Reilly, 2003), until recently little evidence had accumulated to support this view (but see Jones, 1976).

1.2.1. Behavioural evidence for holistic retrieval

Meiser and Bröder (2002) presented evidence consistent with the proposal of holistic retrieval. In this study, participants learnt a series of words that were either presented at an upper or lower location on the screen. At both locations each of the words were presented in either a large or small font size. To assess memory performance, Meiser and Bröder (2002) used a measure of retrieval dependency – the statistical relationship between the retrieval successes of different elements from the same 'event' - to assess whether the probability of retrieving one source correctly (e.g., location) was contingent on the probability of retrieving the other source correctly (i.e., font-size). Meiser and Bröder (2002) observed evidence for dependency in the retrieval of both sources, notably in instances where participants had indicated that they remembered the word being presented during the learning phase, relative to when they indicated merely recognising the word as being presented before (see also Meiser, Sattler, & Weißer, 2008). Meiser et al. (2008) argued that evidence for dependency implied that words and their sources were combined into coherent memory representations, consistent with Tulving's (1983) proposal of the ‘event engram’ as the unit of episodic memory.

However, Starns & Hicks (2005) observed that although two sources (e.g.,, location and font-size) presented with a word appeared to show a degree of dependency, reinstating one source at retrieval did not affect memory for the other source. For instance, participants’ ability to retrieve the correct font-size was unaffected by whether the word appeared in the correct location (i.e., the location the word appeared at during learning) or not at test. Consistent with this, Starns and Hicks (2008) observed that different source features (e.g., colour and location) appeared to be directly bound to the item, but not each other. This suggests that sources associated with an item are retrieved independently of each other, and provides evidence against the proposal by Meiser et al. (2008) suggesting that an item and associated source information are bound into a single memory representation.

Horner and Burgess (2013) noted that the experiments reported by Starns and Hicks (2005, 2008) and Meiser and Bröder (2002) gave primacy to a single element, as the two sources were combined with a single item. However, in order to provide evidence for *or* against Tulving's (1983) proposal, dependency should be assessed for distinct elements (e.g., people, locations, and objects) that are not subordinate to any other element (e.g., colour of the object). Horner and Burgess (2013) had participants learn a series of ‘events’ that consisted of famous people, locations, and common objects. For instance, participants were presented with an event that consisted of *Barack Obama*, *kitchen*, and *hammer*, and asked to imagine all the three elements interacting in a meaningful way as vividly as possible. Horner and Burgess (2013) observed that cueing participants with a single element (e.g., *Barack Obama*) and asking them to retrieve the remaining elements from the same event (e.g., *kitchen* and *hammer*) - with either a cued recall or cued recognition task - led to estimates of retrieval dependency that exceeded what would be expected if each of the elements were retrieved independently of each other. Horner and Burgess (2013) argued that these effects, relative to those observed by Starns and Hicks (2005, 2008), may reflect differences in the binding of distinct event elements compared to the binding of objects and their features, with the former thought to involve the hippocampus, and the latter the perirhinal cortex (Aggleton & Brown, 1999; Davachi, 2006; Diana et al., 2007).

However, when all event elements are encoded simultaneously (as in Horner & Burgess, 2013), it is possible that trial-wise differences in attention can lead to the appearance of dependency at retrieval (e.g., Boywitt & Meiser, 2012; Meiser & Sattler, 2007). For instance, if some events are well attended to at encoding then it is possible that the elements of these events will all be remembered, while those not well attended to will be forgotten in their entirety. Consistent with this, Boywitt and Meiser (2012) observed no evidence for dependency when two sources presented with an item were encoded separately (e.g., item and location, followed by item and font-size). However, Horner and Burgess (2014) demonstrated that dependency was still observed when the individual pairwise associations (e.g., *Barack Obama-kitchen*, *kitchen-hammer*, *hammer-Barack Obama*) that made up an event (i.e., *Barack Obama-kitchen-hammer*) where separately encoded in a ‘closed-loop’ associative structure. Interestingly, no dependency was observed when only two out of the three possible pairwise associations were learnt (e.g., *Barack Obama-kitchen*, *kitchen-hammer*, but not *hammer-Barack Obama*) in a ‘open-loop’ associative structure (Horner & Burgess, 2014), consistent with previous work showing that two separately encoded, but overlapping pairwise associations appear to be retrieved independently of each other (Greeno, James, & DaPolito, 1971; Martin, 1971; Wichawut & Martin, 1971). In contrast to proposals suggesting that holistic retrieval is dependent on the elements being presented together at encoding (Hicks & Starns, 2015), Horner and Burgess (2014) showed that dependency can result from the associative structure of the learnt elements. As such, the individual associations can be encoded independently, but still be retrieved holistically, dependent on the associative structure formed during encoding.

1.2.2. fMRI evidence for holistic reinstatement

There is growing evidence to suggest that neocortical regions involved in the initial acquisition of an event are reactivated during later retrieval (see Danker & Anderson, 2010 for review). Initial evidence for this came from studies observing increased retrieval-related activity in the visual cortex for words associated with pictures during learning (Vaidya, Zhao, Desmond, & Gabrieli, 2002; Wheeler, Petersen, & Buckner, 2000) and increased activity in the auditory cortex during retrievals of words associated with sounds during learning (Wheeler et al., 2000). Consistent with this, more recent studies have shown that the retrieval of distinct stimulus categories correlates with activity in ‘content-specific’ cortical regions (Polyn, Natu, Cohen, & Norman, 2005; Ranganath, Cohen, Dam, & D’Esposito, 2004; Woodruff et al., 2005).

Studies assessing memory for item and context information have shown that the strength of reinstatement for contextual information is related to participants’ subjective experience of recollection (Johnson & Rugg, 2007; Woodruff et al., 2005). Consistent with the role of the hippocampus in recollection, evidence has shown that the degree of reinstatement for word-context associations in the neocortex correlates with activity in the hippocampus during retrieval (Gordon, Rissman, Kiani, & Wagner, 2014; Staresina, Henson, Kriegeskorte, & Alink, 2012). Further, Staresina, Cooper, and Henson, (2013) demonstrated that hippocampal activity during retrieval was associated with a ‘signal flow’ from the region associated with a retrieval cue (e.g., *bagel*) to the region associated with a retrieval target (e.g., *kitchen*); suggesting that the hippocampus mediates the reinstatement of associative information represented in distinct regions of the brain.

Although evidence has accumulated to support the role of the hippocampus in the retrieval of episodic memory, studies assessing the reinstatement of event elements have tended to focus on the retrieval of associated elements that participants are explicitly required to retrieve (but see Kuhl, Johnson, & Chun, 2013). However, a prediction that follows from the proposal that the retrieval of episodic memory is holistic (Tulving, 1983) is that all aspects of an event should be reinstated, even those elements that are incidental to the specific retrieval task (Horner, Bisby, Bush, Lin, & Burgess, 2015). In a recent study, Horner et al. (2015) had participants learn a series of ‘closed-loops’. For instance, for one closed-loop, participants learnt the associations between *Barack Obama-kitchen*, *kitchen-hammer*, and *hammer-Barack Obama*. During retrieval, participants were cued with one element (e.g., Barack Obama) and asked to retrieve an associated element (e.g., hammer). Horner et al. (2015) observed that when participants were required to retrieve *Barack Obama* and *kitchen*, participants also showed increased activity in regions associated with the retrieval of the remaining element, *hammer*, despite no explicit requirement to retrieve this element. The reinstatement of elements incidental to retrieval is consistent with the proposal that episodic memories are retrieved holistically (Tulving, 1983). Importantly, Horner et al. (2015) also observed that the degree of reinstatement for elements incidental to retrieval correlated with hippocampal activity.

1.2.3. Summary

Consistent with Tulving's (1983) proposal that episodic memory retrieval is holistic, evidence has demonstrated that event-based representations are associated with all-or-none retrieval (Horner & Burgess, 2013, 2014). Interestingly, source features associated with an item have been shown to be retrieved independently of each other (Starns & Hicks, 2005, 2008). It is possible that this distinction is related to differences in the associative mechanism supporting the encoding and retrieval of event- and item-based memories in the hippocampus and perirhinal cortex, respectively (Aggleton & Brown, 1999; Davachi, 2006; Mayes et al., 2007). Consistent with this, neuroimaging evidence has demonstrated that the hippocampus not only supports the retrieval of associative information (Gordon et al., 2014; Staresina et al., 2013; Staresina et al., 2012), but also plays a role in the retrieval of associative information that participants are not explicitly required to retrieve (Horner et al., 2015). This is consistent with the proposed role of the hippocampus in episodic recollection (Aggleton & Brown, 1999; Cohen & Eichenbaum, 1993), above and beyond a more general role of the hippocampus in the retrieval of associative memory (Squire & Zola-Morgan, 1991).

1.3. Systems consolidation and the role of the hippocampus in retrieval

Memory consolidation refers to a gradual process of maturation whereby a newly acquired memory is stabilised in the neocortex (Glickman, 1961; McGaugh, 1966). Memory consolidation is typically separated into synaptic and systems consolidation. While synaptic consolidation refers to post-encoding changes on a synaptic and cellular level that are assumed to draw to a close within hours of memory acquisition, systems consolidation refers to a more protracted reorganisation of long-term memory throughout the brain, that may last several days, weeks, months or years (Dudai, Karni, & Born, 2015). For the present purposes I will not discuss the mechanisms of synaptic consolidation (but see Dudai, 2004; Kandel, Dudai, & Mayford, 2014 for reviews), and instead focus on memory consolidation on the systems-level.

1.3.1. Complementary learning systems theory

Theories of systems consolidation have typically focused on the role of the hippocampus in the retrieval of episodic memory. This is in light of evidence that patients with damage to the hippocampus exhibit a temporal gradient to their memory impairment; with memory for remote events being relatively spared compared to more recent events (Schnider, Bassetti, Gutbrod, & Ozdoba, 1995; Squire, Haist, & Shimamura, 1989; but see Cipolotti et al., 2001; Sanders & Warrington, 1971). The most prominent account of system consolidation proposes that long-term retention requires two complementary learning systems; a ‘neocortical processing system' and a 'hippocampal learning system' (McClelland et al., 1995; see also Alvarez & Squire, 1994; Marr, 1971; Squire & Alvarez, 1995).

The complementary learning systems theory holds that the retrieval of an episodic memory involves the reinstatement of distributed activity patterns in the neocortex. In the presence of an event element, the neocortical processing system must be structured in such a way that a single element not only elicits activity for the element itself but also for other elements associated with the same event (McClelland et al., 1995). Activity patterns that correspond to the same event are assumed to be elicited by propagating activity in synaptic connections between cortical neurons involved in representing the individual event elements (McClelland et al., 1995). Each instance of memory retrieval (or re-exposure) is thought to gradually adjust synaptic connections between the relevant cortical neurons until an event representation reaches a relatively stable long-term state in the neocortex (McClelland et al., 1995).

However, representations in the neocortex are fairly dense, with different events exhibiting some degree of representational overlap (McClelland et al., 1995). As such, rapid synaptic adjustments can potentially lead to 'catastrophic interference' (McClelland et al., 1995). Adjustments to synaptic connections between cortical neurons must therefore be gradual; entailing that the adjustment of cortical connections in the neocortex is a protracted process that requires multiple instances of retrieval (or re-exposure) before a stable representation can emerge within the neocortex (but see Tse et al., 2007, 2011). However, interleaved adjustments to cortical connections allows for activity patterns in the neocortical processing system to be adequately stabilised before accommodating new information that may otherwise lead to interference (McClelland & Goddard, 1996).

The initial acquisition and retrieval of an event-based representation therefore necessitates the involvement of a second system. Consistent with previous accounts (Marr, 1971; McNaughton & Morris, 1987; O'Reilly & McClelland, 1994; Treves & Rolls, 1994), the complementary learning systems theory proposes that a hippocampal learning system underpins the initial acquisition of an event-based representation. During retrieval, if activity patterns representing a distinct event element arise in the neocortical system, then this will give rise to pattern completion in the hippocampus, which allows for the reinstatement of any associated elements in the neocortex (McClelland et al., 1995; Marr, 1971; Treves & Rolls, 1994). Each instance of retrieval will result in small adjustments to the cortical connections in the neocortex, to a point where the entire event can be retrieved independently of the hippocampus.

1.3.2. Multiple trace\Transformation theory

The observation that memory for remote events, compared to more recent events, is relatively spared following damage to the hippocampus has been taken as evidence for a time-dependent role for the hippocampus in the retrieval of episodic memory. This is consistent with standard consolidation theories (e.g., Marr, 1971; McClelland et al., 1995; Squire & Alvarez, 1995), such as the complementary learning systems theory. However, it has been argued that there are instances where patients with damage restricted to the hippocampus do not exhibit a temporal gradient to their memory impairment (see Nadel & Moscovitch, 1997 for review), and that assessments commonly used to demonstrate the presence of a gradient may be confounded by semantic memory processing (Moscovitch, Nadel, Winocur, Gilboa, & Rosenbaum, 2006; Winocur, Moscovitch, & Bontempi, 2010), which has been shown to be insensitive to damage restricted to the hippocampus (Cermak & O'Connor, 1983). This has motivated the proposal that representations whose retrieval is dependent on the hippocampus may differ from those retrieved independently of the hippocampus (Nadel & Moscovitch, 1997; Winocur & Moscovitch, 2011).

While the sparse activity patterns utilised by the hippocampus are well-suited for encoding detailed and distinct events, the overlapping activity patterns of the neocortex are thought to be 'biased' towards integrating across multiple events (McClelland et al., 1995; Norman & O'Reilly, 2003). According to the multiple-trace theory (Nadel & Moscovitch, 1997), each instance of retrieval (or re-exposure) leads to the formation of a new representation in the hippocampus. As retrieval is likely to occur across many different contexts, the existence of ‘multiple traces’ in the hippocampus are proposed to allow the neocortex to extract gist-like information that is shared across the varying contexts in which a memory was retrieved (Nadel & Moscovitch, 1997).

The transformation theory (Winocur & Moscovitch, 2011; Winocur et al., 2010) has argued that the neocortical representation that possibly emerges in place of the hippocampus reflects a gist-like or semantic version deprived of the contextual details embedded within the original hippocampal representation. According to this account, differences in memory for remote and recent events observed in patients with hippocampal damage do not necessarily reflect the remoteness of the memory, but may instead be characterised by differences in the ‘quality’ of the underlying memory representation (Winocur & Moscovitch, 2011). This proposal holds that the retrieval of gist-like representations can perhaps be achieved independently of the hippocampus, but for event-based representations that retain their ‘episodic’ quality, retrieval will continue to be dependent on the hippocampus, irrespective of the age of the memory (Winocur & Moscovitch, 2011). The transformation theory does not exclude the possibility that gist-like representations in the neocortex and more distinct ‘episodic’ representations in the hippocampus can coexist. In those instances, the role of the hippocampus in retrieval will be reflected in the specific demands at retrieval and possibly also the relative strength of the neocortical and hippocampal representations (Winocur & Moscovitch, 2011). Support for the transformation theory has come from evidence showing that hippocampal activity does not consistently differ between recent and remote events (Maguire, Henson, Mummery, & Frith, 2001; Ryan et al., 2001; but see Niki & Luo, 2002; Piefke, Weiss, Zilles, Markowitsch, & Fink, 2003), but instead tends to correlate with the quality or vividness of retrieval for both remote and recent events (Addis, Moscovitch, Crawley, & McAndrews, 2004; Gilboa, Winocur, Grady, Hevenor, & Moscovitch, 2004).

1.3.3. Systems consolidation during sleep

Although the debate as to how to conceptualise the process of systems consolidation is still ongoing (e.g., Barry & Maguire, 2019), standard theories of consolidation have tended to agree that sleep may play a role in systems consolidation (e.g., Marr, 1971; McClelland et al., 1995). Notably, it has been proposed that neural activity established during slow-wave sleep may underlie the reinstatement processes thought to be the core mechanism of systems consolidation (Born, 2010; Diekelmann & Born, 2010).

Slow-wave sleep is characterised by three distinct neuronal oscillations (i.e., slow-oscillations, spindles, and sharp-wave ripples) that are thought to provide the temporal structure for systems consolidation during sleep (Buzsáki, 1996). Slow-oscillations reflect alternating states of hyper- and depolarising neuronal activity at ~.75 Hz (Achermann & Borbély, 1997), and represent 'travelling waves' of activity that emerge in the neocortex (Massimini, Huber, Ferrarelli, Hill, & Tononi, 2004) and extend into the hippocampus (Nir et al., 2011). During hyperpolarising states, slow-oscillations have been related to the onset of sleep spindles (Mölle, Marshall, Gais, & Born, 2002). Spindles involve short bursts of neural activity at ~10-15 Hz (Staresina et al., 2015) that are thought to support hippocampal-neocortical 'dialogue' during sleep; due to their precise temporal relationship with slow oscillations and sharp-wave ripples (Siapas & Wilson, 1998; Sirota & Buzsáki, 2005). Sharp-wave ripples refer to high-frequency bursts of activity at ~80-140 Hz in the hippocampus (Axmacher, Elger, & Fell, 2008; Staba, Wilson, Bragin, & Fried, 2002) that are clustered within spindle 'troughs' (Staresina et al., 2015). Evidence from rodent studies has linked sharp-wave ripples to the reactivation of memory during sleep (Dupret, O'Neill, Pleydell-Bouverie, & Csicsvari, 2010; Wilson & McNaughton, 1994).

The coordinated activity between slow-oscillations, sleep spindles and sharp-wave ripples has been proposed to underlie systems consolidation during sleep (Buzsáki, 1996; Sirota, Csicsvari, Buhl, & Buzsáki, 2003). Consistent with this, memory-related improvements have been shown to be associated with variations in slow-oscillations and spindle properties during sleep (Bergmann, Mölle, Diedrichs, Born, & Siebner, 2012; Gais, Mölle, Helms, & Born, 2002; Marshall, Helgadóttir, Mölle, & Born, 2006; Mednick et al., 2013; Schabus et al., 2004). Furthermore, neuroimaging results have shown that reactivation of recently formed memories during sleep predicts subsequent memory performance (Deuker et al., 2013; Peigneux et al., 2004), providing support for the proposal that memory reactivation during sleep may be involved in the reinstatement processes underlying systems consolidation (Born, 2010; Diekelmann & Born, 2010)

Neuroimaging studies have demonstrated that neocortical activity associated with successful memory retrieval increases following sleep (Gais et al., 2007; Takashima et al., 2009), providing support for the proposed role of sleep in memory consolidation in the neocortex (Born, 2010; Diekelmann & Born, 2010). Consistent with standard theories of consolidation (Alvarez & Squire, 1994; McClelland et al., 1995), it has also been shown that the duration of slow-wave sleep during a nap session correlates negatively with hippocampal activity during memory retrieval one day and 90 days after encoding (Takashima et al., 2006). This finding is consistent with the proposal that systems consolidation during sleep may be related to the formation of a neocortical representation that can be retrieved independently of the hippocampus (Born & Wilhelm, 2012).

Interestingly, a study by Sterpenich et al. (2009) showed that performance for memories judged as recollected 6 months after learning was associated with activity in the neocortex, but not the hippocampus, in participants' who slept following encoding. This pattern of activity was not seen in participants who were sleep-deprived on the first night following encoding, suggesting that post-encoding sleep may play a distinct role in systems consolidation (Sterpenich et al., 2009). Consistent with this proposal, evidence suggests that sleep, relative to an equivalent period of daytime wakefulness, supports participants' ability to gain insight into abstract rules in a sequence learning task (Wagner, Gais, Haider, Verleger, & Born, 2004) and participants' ability to infer the relationship between pairs of items embedded in an hierarchical structure on a relational memory task (Ellenbogen, Hu, Payne, Titone, & Walker, 2007). This has been taken as evidence to suggest that sleep may play an active role in the generalisation processes required to extract commonalities across related or overlapping events (Diekelmann & Born, 2010; Lewis & Durrant, 2011; Stickgold & Walker, 2013).

Other studies have noted that decreases in hippocampal activity following sleep appear to be restricted to the posterior hippocampus, with activity in the anterior hippocampus during memory retrieval remaining relatively stable following a one day delay that includes a night of sleep (Ritchey, Montchal, Yonelinas, & Ranganath, 2015; Takashima et al., 2009). It is possible that this difference may relate to evidence suggesting a functional distinction between the posterior and anterior axis of the hippocampus (Collin, Milivojevic, & Doeller, 2015; Poppenk & Moscovitch, 2011). It has been argued that the posterior hippocampus may play a greater role in the encoding and retrieval of contextual details, relative to the anterior hippocampus (Poppenk, Evensmoen, Moscovitch, & Nadel, 2013). Decreases in retrieval-related activity in the posterior hippocampus may therefore reflect the forgetting of contextual information (Ritchey et al., 2015). As such, it is possible that these findings are more consistent with the proposal that systems consolidation may be associated with the formation of a gist-like representation in the neocortex, that is devoid of the 'episodic' quality of the original, hippocampal representation (Winocur & Moscovitch, 2011).

1.3.4. Summary

Although the hippocampus has been implicated in the encoding and retrieval of episodic memory (Aggleton & Brown, 1999; Cohen & Eichenbaum, 1993), it is less clear whether the hippocampus continues to support retrieval after a period of consolidation. Standard theories of consolidation have suggested that systems consolidation is associated with the formation of a neocortical representation that can be retrieved independently of the hippocampus (Marr, 1971; McClelland et al., 1995). However, the transformation theory has argued that the neocortical representation that potentially emerges in place of a hippocampal representation must reflect a gist-like version of the original hippocampal representation (Winocur & Moscovitch, 2011; see also Nadel & Moscovitch, 1997). Thus, retrieval independent of the hippocampus will reinstate a memory that is devoid of the contextual details associated with the original experience. Yet, the transformation account does not exclude the possibility that neocortical and hippocampal representations may coexist, and as such, retrieval will continue to be dependent on the hippocampus for representations that retain their ‘episodic’ quality.

1.4. Forgetting as a function of interference, inhibition, and decay

1.4.1. Forgetting due to interference and inhibition

Following memory acquisition, representations are prone to forgetting. Historically, the prevailing assumption of forgetting was that encoding of new information can disrupt or be disrupted by an existing memory representation (see Wixted, 2004 for review). For instance, Jenkins and Dallenbach (1924) demonstrated that retention for nonsense syllables was increased in participants who slept between encoding and retrieval, relative to those who remained awake. As participants who remained awake were expected to be exposed to more interfering material following encoding, the finding that sleep reduced forgetting was taken to suggest that new memory formation throughout the day, in participants who remained awake, retroactively interfered with the critical test information. However, Underwood (1957) later argued that much of forgetting observed in previous experiments was accounted for by the material participants learnt *prior*, rather than *after*, the critical test information, suggesting that retention was at least as susceptible to proactive as retroactive interference (see Postman, 1971 for review).

More recently it has been argued that forgetting is not always a consequence of new or past learning *per se*, but can also be the result of inhibitory processes established during retrieval to overcome interference (Anderson, 2003). This proposal holds that when a situation demands the selective retrieval of a memory, inhibitory processes present at retrieval will lead to long-lasting decreases in retention for other memories that may compete or interfere with the 'target' memory (Anderson, 2003). This proposal is motivated by evidence presented by Anderson, Bjork, and Bjork (1994). Anderson et al. (1994) had participants learn six category exemplars of, for instance, fruits and asked participants to repeatedly retrieve three of the six exemplars in response to a cue exemplar stem (e.g., *fruit: or_* for the exemplar *orange*). Anderson et al. (1994) observed that not only did performance for the retrieved exemplars increase on a later memory test, but memory for non-retrieved exemplars (e.g., *banana*) decreased relative to a baseline category of exemplars. This phenomenon is referred to as 'retrieval-induced forgetting' (Anderson et al., 1994). Retrieval-induced forgetting appears to emerge as an effect of interference at retrieval, with greater decreases in memory performance seen for exemplars that show a greater amount of relatedness to the target memory (Anderson et al., 1994). This effect is not only seen for exemplars within the same category, but extends to exemplars that are from different categories but nevertheless relate to a 'target' memory (Anderson & Spellman, 1995), suggesting that the effect of repeated retrieval can generalise to any representation that interferes with a target memory.

The idea that retrieval-induced forgetting is related to inhibitory processes at retrieval is supported by neuroimaging work. For instance, Kuhl, Dudukovic, Kahn, and Wagner (2007) had participants learn a series of overlapping word pairs (e.g., *attic-dust* and *attic-junk*). During retrieval practice, participants were repeatedly presented with a cue word and the first letter of the associated word required to be retrieved (e.g., *attic-d___* for the associate *dust*).

Kuhl et al. (2007) observed that activity in the prefrontal cortex – a region associated with inhibitory control processes (see Miller & Cohen, 2001 for review) – decreased across the first and final retrieval practice trial. The reduction in activity in the prefrontal cortex correlated with decreases in memory for the ‘unpractised’ association (i.e., *attic-junk*) on a final memory test, suggesting that inhibitory demands during retrieval practice decreased as the unpractised association was forgotten (see also Wimber, Rutschmann, Greenlee, & Bäuml, 2008). Further evidence for this was presented by Wimber, Alink, Charest, Kriegeskorte, and Anderson (2015). Wimber et al. (2015) had participants learn a series of overlapping associations (e.g., *sand-Marilyn Monroe*, *sand-hat*), and then asked participants to selectively retrieve one of the associative items (e.g., *Marilyn Monroe*), but not the other item (e.g., *hat*) in response to a retrieval cue (e.g., *sand*). During repeated retrievals, Wimber et al. (2015) observed that cortical activity patterns became more dissimilar to activity patterns associated with the unpractised item. This was taken as evidence that retrieval practice led to the ‘suppression’ of cortical representations of the unpractised item (Wimber et al., 2015). Consistent with this, the extent of this cortical suppression effect was shown to correlate with the forgetting of a non-practised item (e.g., *hat*) (but see Potter, Huszar, & Huber, 2018). This suppression effect was also shown to correlate with activity in the prefrontal cortex during retrieval practice, consistent with the proposal that inhibitory processes can lead to forgetting when two or more memory representations interfere with each other at retrieval (Anderson, 2003).

1.4.2. Forgetting as a function of decay

Interference from ongoing learning is not thought to be the primary mechanism for forgetting of memory representations supported by the hippocampus (Sadeh, Ozubko, Winocur, & Moscovitch, 2014, 2016). As representations in the neocortex are overlapping in nature, information represented in the neocortex is likely to be susceptible to interference (McClelland et al., 1995; Norman & O’Reilly, 2003). However, the sparse activity patterns used by the hippocampus are thought to support pattern separation (Bakker et al., 2008; Berron et al., 2016). As such, memory representations in the hippocampus are less likely to interfere with each other (Yassa & Stark, 2011). Consistent with this, activity in the hippocampus during the repeated encoding of an A-C pair (e.g., *watch-pipe*) has been shown to support retention of a previously learnt, and overlapping A-B pair (e.g., *watch-sink*) (Kuhl, Shah, Dubrow, & Wagner, 2010; see also Chanales, Dudukovic, Richter, & Kuhl, 2019), suggesting that activity in the hippocampus minimises the likelihood of forgetting due to interference during encoding.

Sadeh, Ozubko, Winocur, and Moscovitch (2014) proposed that differences in the mechanisms of forgetting may be reflected in the representations and/or retrieval processes that support memory. Specifically, Sadeh et al. (2014) suggested that familiarity - a process thought to be supported by the perirhinal cortex (Diana et al., 2007; Mayes et al., 2007) – is sensitive to forgetting via interference. However, because pattern separation in the hippocampus ensures that memory representations are less likely to interfere with each other (Yassa & Stark, 2011), recollection – a process thought to be supported by the hippocampus (Diana et al., 2007; Mayes et al., 2007; Yonelinas, 2002) – is more likely to be forgotten via decay (Sadeh et al., 2014). Some support for this proposal comes from evidence showing that while memory for words judged as 'remembered' decreases between encoding and retrieval after a one week-delay, recognition for words judged as 'known' does not decrease as a function of delay (Gardiner & Java, 1991; but see Hockley & Consoli, 1999). However, Sadeh, Ozubko, Winocur, and Moscovitch (2016) provided more direct evidence for this proposal by manipulating both the effect of interference and delay on retention for words judged as either recollected or familiar. Sadeh et al. (2016) observed that while memory for words judged as familiar decreased predominately as an effect of interference (i.e., presenting participants with additional words following encoding), memory for words judged as recollected decreased as a function of the length of the delay between encoding and retrieval, irrespective of the degree of interference (consistent with forgetting via decay).

The proposal by Sadeh et al. (2014) is motivated by previous work suggesting that the forgetting of hippocampal representations is the result of decay (Frankland, Kohler, & Josselyn, 2013; Hardt, Nader, & Nadel, 2013); possibly due to ongoing neurogenesis in the hippocampus (Frankland et al., 2013) or more active regulatory processes during sleep that 'remove' hippocampal representations based on their behavioural relevance or representational strength (Hardt et al., 2013). The proposal that forgetting of hippocampal representations is the result of decay has some parallels with theories of systems consolidation (e.g., McClelland et al., 1995; Squire & Alvarez, 1995) in that it assumes that decay involves the gradual deterioration of a hippocampal representation, leaving only the neocortical components of the memory behind. Prior to consolidation, the decay of a hippocampal representation can leave neocortical representations prone to interference and lead to forgetting (Frankland et al., 2013; Hardt et al., 2013). However, for neocortical representations that have been sufficiently strengthened via consolidation, the consequence of decay in the hippocampus is thought to be limited to the loss of contextual information associated with a particular event (Hardt et al., 2013); consistent with the proposal that systems consolidation may be associated with a formation of a gist-like representation in the neocortex (Winocur & Moscovitch, 2011). Thus, forgetting will reflect an interaction between two opposing processes, memory consolidation and memory decay (Frankland et al., 2013).

1.4.3. Summary

Research on forgetting has tended to focus on the effect of interference during learning (see Wixted, 2004 for review). Although it is accepted that some of forgetting will be due to interference, it has been argued that forgetting is not always due to interference during learning, but can also reflect inhibitory processes that are established to overcome interference at retrieval (Anderson, 2003). However, more recently it has been suggested that some of forgetting may also be due to memory decay (Frankland et al., 2013; Hardt et al., 2013; Sadeh et al., 2016). This proposal draws on our neuroscientific understanding of the anatomical characteristics that distinguish medial temporal lobe regions, such as the hippocampus, from extra-hippocampal regions and the neocortex more generally. Due to the presence of pattern separation in the hippocampus, the forgetting of representations, and/or retrieval processes, supported by the hippocampus is not thought to occur via interference. Instead, it has been proposed that forgetting of a hippocampal representation may be due to the decay of the underlying memory trace itself.

1.5. Thesis overview

Research has accumulated to provide support for Tulving's (1983) proposal that episodic memories are retrieved in a holistic manner. However, whereas Tulving (1983) emphasised the spatio-temporal context as the core feature of a distinct episodic memory, recent work has demonstrated that coherent event-based representations can be formed via the associations between the individual elements themselves; in the absence of a distinct spatial-temporal context (Horner et al., 2015; Horner & Burgess, 2014). Horner and Burgess (2014) had participants separately encode three pairwise associations (e.g., *Barack Obama-kitchen*, *kitchen-hammer*, *hammer-Barack Obama*) in a closed-loop, and observed that retrieval dependency was seen for all elements of this separately encoded event. This dependency was similar to that seen for events where all three elements were encoded on a single trial (Horner & Burgess, 2014), suggesting that the explicit encoding of all possible pairwise associations formed a coherent event representation similar to that formed in a single spatio-temporal context. However, no dependency was seen when only two out of the three pairwise associations (e.g., *Barack Obama-kitchen, kitchen-hammer,* but not *hammer-Barack Obama*) where encoded in an open-loop (Horner & Burgess, 2014). This suggests that the all-to-all structure formed when all possible pairwise associations are encoded in a closed-loop may be a necessary boundary condition for the emergence of a coherent event-based representation that can be retrieved holistically (Horner et al., 2015). Consistent with this, Horner et al. (2015) provided fMRI evidence to demonstrate that the retrieval of closed-loops was associated with neocortical reinstatement of all event elements, even those incidental to the specific retrieval task. For instance, cueing participants with *kitchen* and asking them to retrieve *hammer* was accompanied with increases in activity in neocortical regions associated with the retrieval of *Barack Obama*, despite this element not being on-screen nor being required to be retrieved. This ‘incidental’ reinstatement effect was seen for ‘events’ learnt in a closed-loop, but not those learnt in an open-loop (Horner et al., 2015).

In the experimental chapters that follow, I apply the paradigm developed in Horner and Burgess (2014). This paradigm allows for building ‘events’ that are matched in the number of elements (i.e., people, locations, and objects), but differ in their coherence at retrieval depending on the associative structure formed at encoding. Thus, we can distinguish between closed-loops that are retrieved in a holistic or coherent manner, and open-loops that are retrieved in an independent or less coherent way; and ask how these differences in the associative structure formed during encoding affect retrieval, not only immediately after encoding but also following a period of consolidation and forgetting.

1.5.1. Chapter 2

Brady, Konkle, Alvarez, and Oliva (2013) recently used retrieval dependency to assess the statistical relationship between the retrieval of two features associated with an item. Brady et al. (2013) saw evidence for dependency immediately after encoding, but dependency decreased over time. This decrease in dependency is consistent with the idea that features associated with an item are forgotten independently of each other (Brady et al., 2013). However, whereas item-based representations are thought to be supported by the perirhinal cortex (Aggleton & Brown, 1999; Diana et al., 2007), the hippocampus has been shown to play an important role in the encoding and retrieval of event-based representations (Aggleton

& Brown, 1999; Cohen & Eichenbaum, 1993; Eichenbaum et al., 2007; Mayes et al., 2007). It has been proposed that forgetting processes may differ between these two regions, with representations/retrieval processes supported by the perirhinal cortex being more susceptible to forgetting via interference, and representations/retrieval processes supported by the hippocampus being more likely to be forgotten via decay (Sadeh et al., 2014).

In Chapter 2, I asked whether the elements that make up an event-based representation are also forgotten independently (as in Brady et al., 2013), or whether event-based representations are instead more likely to be forgotten in an all-or-none/holistic manner. In Chapter 2, I present four experiments that assess retrieval accuracy and dependency for closed- and open-loops immediately after encoding and following a delay. I assessed whether decreases in retrieval accuracy between immediate and delayed retrieval are accompanied by decreases in dependency. This would imply that closed-loops fragment as a function of forgetting. However, if dependency remains consistent over time, then this would imply that closed-loops are instead more likely to be forgotten in an all-or-none manner. In Experiment 1, I assessed retrieval accuracy and retrieval dependency immediately after encoding and following a 12-hour delay. I also included a sleep/wake manipulation, such that half of the participants slept between encoding and retrieval, and the other half remained awake between encoding and retrieval. This is because previous work has suggested that systems consolidation during sleep may play an active role in the integration of overlapping memory representations (Diekelmann & Born, 2010; Lau, Tucker, & Fishbein, 2010), and as such I was also interested in assessing whether post-encoding sleep may be associated with increases in dependency for open-loops formed of two overlapping pairwise associations. In Experiments 2-4, I assessed retrieval accuracy and retrieval dependency immediately after encoding (Experiments 2 and 4) and following a one-week delay (Experiment 2-4) in order to increase the amount of forgetting between the two retrieval time points.

1.5.2. Chapter 3

While Chapter 2 principally focussed on the forgetting of closed-loops, in Chapter 3 I turned my attention to the forgetting and consolidation of open-loops. Previous work has shown that the encoding of two out of the three possible pairwise associations in an open-loop does not lead to dependency at retrieval (Horner & Burgess, 2014), suggesting that the associations encoded as open-loops are retrieved independently. (Horner & Burgess, 2014). However, systems consolidation during sleep has been proposed to strengthen the association between overlapping memory representations (Diekelmann & Born, 2010; Lewis & Durrant, 2011). Consistent with this proposal, a short nap following encoding has been shown to improve participants' inference performance on non-encoded B-C pairs, following the encoding of overlapping A-B, A-C pairs (Lau, Tucker, & Fishbein, 2010) (analogous to the open-loop associative structure used here). Although dependency is not typically seen for open-loops immediately following encoding, systems consolidation during sleep may support mnemonic integration for open-loops, such that dependency is seen following sleep.

I assessed retrieval accuracy and dependency for open-loops immediately following encoding and after a period of sleep. As memory reactivation is proposed to be the core mechanism of systems consolidation (McClelland et al., 1995; Squire & Alvarez, 1995), I used a technique called targeted memory reactivation (TMR) (Rudoy, Voss, Westerberg, & Paller, 2009) to induce memory reactivation of open-loops during sleep, and assessed if retrieval accuracy and retrieval dependency differed depending on whether open-loops were subject to TMR during sleep or not (rather than sleep more generally as in Chapter 2). Although TMR during sleep has been shown to support retention for single associative memories (see Cellini & Capuozzo, 2018 for review), more recent work has demonstrated that TMR for two overlapping associations can, in some instances, lead to forgetting (Antony, Cheng, Brooks, Paller, & Norman, 2018; Oyarzún, Morís, Luque, de Diego-Balaguer, & Fuentemilla, 2017). Thus, in Chapter 3, I was also interested in assessing whether TMR of open-loops during sleep leads to differences in retrieval accuracy for overlapping associations, possibly as a function of proactive interference at encoding.

1.5.3. Chapter 4

In Chapter 4, I returned to the proposal that closed-loops perhaps tend to be forgotten in an all-or-none manner. Horner et al. (2015) previously used fMRI to demonstrate that closed-loop elements are reinstated during retrieval, including those incidental to the specific retrieval task. This is consistent with the proposal that event-based memories are associated with holistic retrieval (Tulving, 1983). In Chapter 4, I used fMRI to assess whether closed-loops continue to be reinstated in a holistic manner following a period of forgetting and consolidation.

Horner et al. (2015) also demonstrated that the reinstatement of elements incidental to retrieval correlated with activity in the hippocampus for closed-loops retrieved immediately after encoding, providing support for the proposed role of the hippocampus in the holistic retrieval of episodic memory (Aggleton & Brown, 1999; Cohen & Eichenbaum, 1993; Eichenbaum et al., 2007; Mayes et al., 2007). However, theories of systems consolidation make different predictions about the role of the hippocampus following a period of consolidation. While standard theories of consolidation (e.g., McClelland et al., 1995; Squire & Alvarez, 1995) suggest that the role of the hippocampus in retrieval decreases following a period of consolidation, the transformation theory (Winocur & Moscovitch, 2011) argues that the hippocampus will continue to be involved in the retrieval of event-based representations that retain their ‘episodic’ quality, regardless of the interval between memory acquisition and retrieval. In light of these differing predictions, I also assessed whether activity in the hippocampus continues to correlate with the neocortical reinstatement of elements incidental to retrieval, regardless of whether a closed-loop is retrieved immediately after encoding or following a period of forgetting and consolidation.

2. All-or-none forgetting of closed-loops

The experiments presented in Chapter 2 have previously been published as: **Joensen, B. J., Gaskell, M.G., & Horner, A. J. (2019). United we fall: All-or-none forgetting of complex episodic events. *Journal of Experimental Psychology General*. *149*(2), 230-248**. The text of Joensen et al. (2019) has been altered for the present chapter to ensure greater continuity between thesis chapters but is predominately the same as in the listed publication.

2.1. Abstract

Research suggests that item-based representations are forgotten in a fragmented manner. However, we do not know how more complex event-based representations are forgotten. Here I assessed whether retrieval accuracy and dependency for event-based representations decreases over a period of forgetting. Across four experiments, retrieval dependency was seen both immediately after learning and following a 12-hour and one-week delay. Further, the amount of retrieval dependency after a delay was greater than that predicted by a model of independent forgetting. This dependency was only seen for closed-loops; where all pairwise associations between locations, people and objects were encoded. When open-loops were learnt; where only two out of the three possible associations were encoded, no dependency was seen immediately after learning or after a delay. I also observed evidence for higher retention rates for closed-loops than open-loops, but only when participants learnt both closed- *and* open-loops, rather than when participants learnt either closed- *or* open-loops. These findings suggest that event-based representations are not only retrieved, but also forgotten, in an all-or-none manner.

2.2. Introduction

Early research on forgetting was dominated by a theoretical debate concerning whether forgetting occurs as a function of interference; where overlapping memory traces disrupt one another, or decay; where memory traces decay over time (see Wixted, 2004 for review). Evidence for greater forgetting of nonsense syllables when participants remained awake, relative to when participants slept, between encoding and retrieval were taken as evidence for interference accounts (Jenkins & Dallenbach, 1924), as decay was thought to lead to equal rates of forgetting across wake and sleep. As more interfering material would be encoded in the awake, relative to sleep, condition following learning, the greater rate of forgetting for participants who remained awake was taken as evidence for retroactive interference. Further evidence for the interference account was provided by Underwood (1957). However, Underwood (1957) showed that the more material learnt *prior* to the critical test information, the greater the subsequent forgetting. Thus, in contrast to the findings by Jenkins and Dallenbach (1924), forgetting appeared to occur predominantly as a function of proactive interference (see Postman, 1971 for review).

More recently, the idea that forgetting is principally a function of interference, and proactive interference in particular, has been questioned. For instance, Wixted (2004) proposed that forgetting is predominantly a result of retroactive interference from mental activity or new memory formation more generally, as opposed to the specific interference that occurs from learning related/overlapping materials. Further, the concept of memory decay has also been revived (Frankland et al., 2013; Hardt et al., 2013). This account holds that some of forgetting will be due to the deterioration of the biological substrates of a memory representation itself. Importantly, both proposals draw on our neuroscientific understanding of forgetting and the concept of consolidation; where new memory representations are thought to stabilise over time, becoming less susceptible to interference and/or decay (see Dudai et al., 2015 for review).

Although, dual-process memory models have proposed different rates of forgetting dependent on the type of representations (Brainerd & Reyna, 2002; Reyna & Brainerd, 1995) or that different representations are more likely to be forgotten via decay or interference (Sadeh et al., 2014), relatively little research has focussed on whether mnemonic representations change as a function of forgetting. Brady et al. (2013) recently used forgetting rates to infer the representational structure of item-based memory representations. Brady et al. (2013) observed different forgetting rates for separate aspects of an object (i.e., colour and state), with the colour of an object being forgotten more rapidly than the state of the object (i.e., its pose or configuration of parts). These results suggest that item-based representations fragment over time, with some aspects of the memory trace being forgotten more rapidly than others. Brady et al. (2013) also assessed retrieval dependency. The presence of dependency has previously been used to infer the coherence of the underlying memory representation (Horner & Burgess, 2013, 2014). Consistent with previous research (Meiser & Bröder, 2002), Brady et al. (2013) saw evidence for dependency after initial encoding, but importantly dependency decreased over time. This decrease in dependency is consistent with a fragmentation of the memory trace as a function of forgetting.

Here, I asked whether more complex event-based representations also fragment over time. Whereas item-based memories can be supported by medial temporal lobe regions outside of the hippocampus, such as the perirhinal cortex (Aggleton & Brown, 1999; Diana et al., 2007), the hippocampus is thought to be critical for the encoding and retrieval of event-based representations (Cohen & Eichenbaum, 1993; Scoville & Milner, 1957; Squire & Zola-Morgan, 1991; Vargha-Khadem, 1997). Thus, these two types of mnemonic representations may rely on distinct regions of the medial temporal lobe (but see Song, Wixted, Hopkins, & Squire, 2011; Wais, Wixted, Hopkins, & Squire, 2006). This point is critical given the recent proposal that forgetting processes may differ between these two regions (Sadeh et al., 2014). Research has suggested that whereas recollection – a process thought to be supported by the hippocampus (Diana et al., 2007; Mayes et al., 2007) – is more likely to be forgotten via decay processes, familiarity – a process thought to be supported by the perirhinal cortex (Diana et al., 2007; Mayes et al., 2007) – is more likely to be forgotten via interference (Sadeh et al., 2016). Given these dissociations, it is possible that event-based representations, supported by the hippocampus, do not undergo the same fragmentation process as seen in Brady et al. (2013), but are instead forgotten in a more all-or-none manner.

It has previously been shown that the encoding of overlapping pairwise associations can result in retrieval dependency (Horner & Burgess, 2014). For example, if a participant learns the associations between *kitchen* and *hammer*, *kitchen* and *Barack Obama*, and *hammer* and *Barack Obama* in a closed-loop across three separate encoding trials, retrieval dependency is seen for all constituent elements of this separately encoded 'event' – i.e., if a participant is cued with *kitchen* and successfully retrieves *Barack Obama*, the participant is more likely to also successfully retrieve *hammer* when cued with *kitchen* on a separate retrieval trial. This dependency is similar to that seen when all three elements are encoded on a single trial (Horner & Burgess, 2013, 2014), suggesting that encoding all three pairwise associations forms a coherent memory representation similar in nature to that formed in a single spatiotemporal context. Horner et al. (2015) also provided fMRI evidence to suggest that these elements are bound into coherent 'event engrams' in the hippocampus, allowing for the subsequent retrieval of all event elements. The retrieval of all event elements is consistent with the idea that recollection (as opposed to familiarity) is associated with all-or-none retrieval (Norman & O'Reilly, 2003), and provides evidence for all-or-none retrieval occurring via pattern completion (Gardner-Medwin, 1976; Hopfield, 1982; Marr, 1971; McClelland, McNaughton, & O'Reilly, 1995; Treves & Rolls, 1992; see Horner & Doeller, 2017; Hunsaker & Kesner, 2013 for review).

Here, I define ‘forgetting’ as decreases in retrieval accuracy between two retrieval time points, and remain theoretically agnostic as to whether forgetting is a result of a decreased accessibility for intact memory traces, or a loss of the underlying trace itself (Tulving & Pearlstone, 1966). Although any measure of forgetting will inevitably be derived from differences in the proportion of memories retained between two retrieval time points, here I asked, as closed-loops are forgotten (as measure by retrieval accuracy), does retrieval dependency decrease. This would imply that the underlying memory traces are fragmenting over time (see Figure 1C). However, if forgetting occurs, but dependency is consistent over time, then this would imply that coherent event-based representations are instead forgotten in an all-or-none manner (see Figure 1B), with event-based representations being more likely to either be retained or forgotten in their entirety.

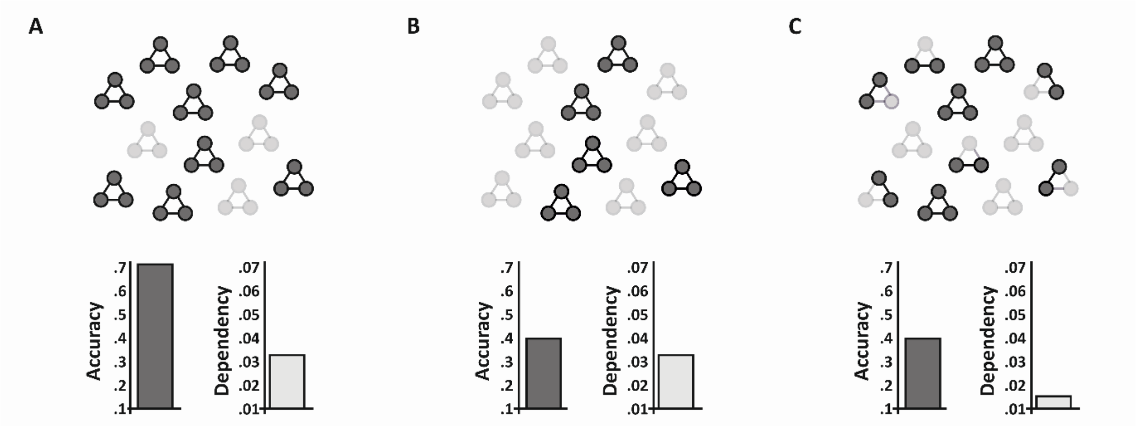

Figure 1 Illustration of retrieval accuracy and dependency for three-element closed-loops immediately after encoding and after all-or-none or independent forgetting has occurred. **(A)** After encoding. It is assumed that some closed-loops are either not encoded or are forgotten between encoding and immediate test (represented by transparent events). Retrieval is all-or-none for remembered events. All-or-none retrieval is reflected in values of dependency significantly greater than 0 (see Methods for how dependency is assessed). **(B)** After all-or-none forgetting. Closed-loops are forgotten in an all-or-none manner. Despite decreases in retrieval accuracy, due to forgetting, dependency does not decrease relative to **A**. **(C)** After independent forgetting. Individual associations are remembered and/or forgotten within the same closed-loop. Dependency decreases relative to **A**, despite the same decrease in accuracy as in **B**.

I used a design similar to that developed in Horner and Burgess (2014). Across all four experiments, participants encoded a series of multi-element ‘triplets’ (Figure 2). Each triplet consisted of three elements (famous people, locations and objects). Triplets were ‘built up’ over two/three separate, spaced, encoding trials. Each trial consisted of the presentation of one of the three possible pairwise associations. This allows for the formation of triplets with different structures of overlapping pairs; ‘closed-loops’; where all the pairwise associations are encoded (e.g., *kitchen-hammer*, *kitchen-Barack Obama*, *hammer-Barack Obama*), or ‘open-loops’; where only two out of the three possible pairwise associations are encoded (e.g., *kitchen-hammer, kitchen-Barack Obama,* but not *hammer-Barack Obama*).

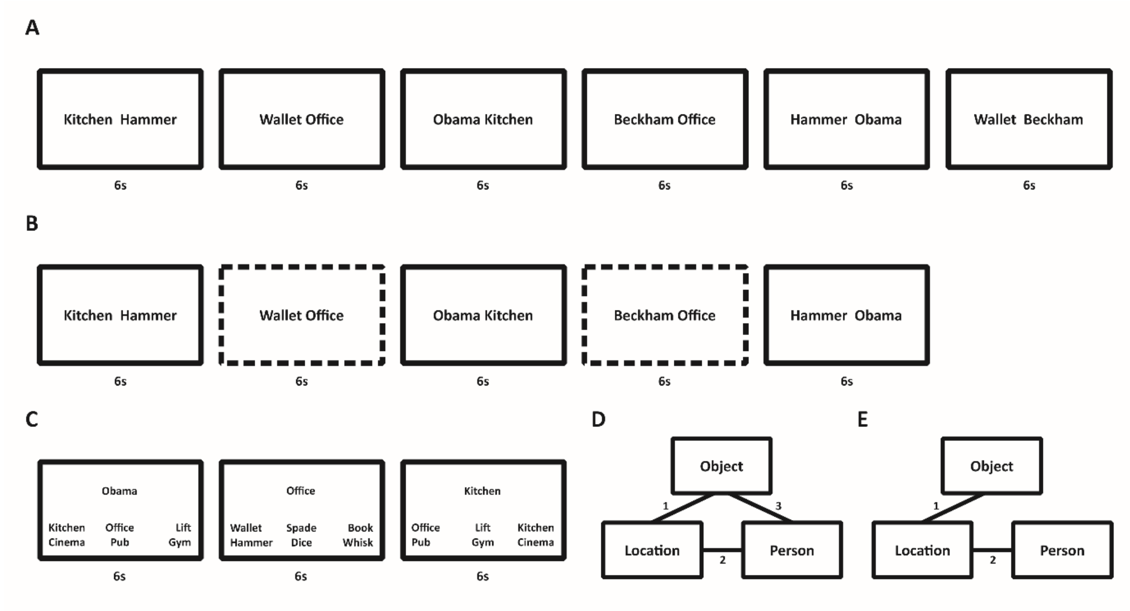

Figure 2 Experimental design. **(A, B)** Encoding. Participants saw multiple pairwise associations. They imagined each association ‘interacting in a meaningful way as vividly as possible’ for 6s. Each association was preceded by a 500-ms fixation and followed by a 500-ms blank screen **(A)** Experiments 1-2. Participants encoded two or three overlapping pairwise associations depending on whether they were allocated to the between-subject open- vs closed-loop condition, respectively. In the open-loop condition, participants did not encode the third and final association (e.g., *hammer-Obama* and *wallet-Beckham*) (see **E**). **(B)** Experiments 3-4. Participants encoded open- *and* closed-loop pairwise associations in an intermixed manner. Solid and dotted lines were not presented but highlight closed-(solid lines) and open-loops (dotted lines). **(C)** Retrieval. Participants were presented with a single cue and required to retrieve one of the other elements from the same event from among five foils (elements of the same type from other events) in 6s. Each cued-recognition trial was preceded by a 500-ms fixation and followed by a 500-ms blank screen. **(D)** The associative structure of closed-loops with example encoding order for three pairwise associations (numbers 1-3). **(D)** The associative structure of open-loops with example encoding order for the two pairwise associations (numbers 1-2). The third and final association (i.e., person-object in this example) is not shown to the participants.

It has previously been shown that dependency is seen for closed-loops (and three element triplets learnt on a single encoding trial), but not for open-loops (Horner et al., 2015; Horner & Burgess, 2013, 2014). The associative structure formed for closed-loops is therefore similar in nature to the coherent memory representations formed on single encoding trials (Horner & Burgess, 2014). Given these findings, closed-loops are referred to as 'events', but note that they are not single spatiotemporal events as typically defined (Tulving, 1983). Dependency is not seen for open-loops, and as such the open-loop condition serves as a control condition where dependency is not expected, even when retrieval shortly follows encoding. The inclusion of the open-loop condition also allows for assessing the further possibility that overlapping associations may undergo a process of integration over time, such that open-loops might show dependency after a delay. This is in light of research showing that the ability to infer the relationship between non-encoded B-C pairs (after encoding A-B and A-C pairs) increases following a short nap (Lau et al., 2010). Thus, sleep may play a role in generalising across related, but independently encoded, information (Diekelmann & Born, 2010); as is the case for open-loops.

At retrieval, encoded associations were tested in both directions (e.g., cue *kitchen*, retrieve *Barack Obama*; cue *Barack Obama*, retrieve *Kitchen*) using cued six-alternative forced-choice. In Experiment 1, retrieval was tested immediately and following a 12-hour delay. At immediate retrieval only encoded associations from half of the triplets were tested, while at delayed retrieval encoded associations from all triplets were tested. Forgetting was presumed to have occurred after 12 hours relative to the immediate condition. Given the well-established finding that sleep decreases forgetting (see Diekelmann & Born, 2010 for review), the extent of forgetting was also manipulated by encoding in the morning or evening, such that half the participants remained awake between encoding and retrieval and half slept. The sleep manipulation also allows for assessing whether sleep may play a role in integrating pairwise associations encoded as open-loops (e.g., Lau et al., 2010).

2.3. Experiment 1

Experiment 1 assessed retrieval accuracy and dependency immediately after encoding and after a 12-hour delay. The time of encoding was manipulated, such that half of the participants encoded pairwise associations in the morning and half of the participants encoded in the evening. This meant that half of the participants slept between encoding and retrieval (i.e., sleep condition), and half remained awake between encoding and retrieval (i.e., awake condition). The inclusion of the sleep manipulation was twofold: (1) to vary the amount of forgetting whilst controlling for the interval between initial encoding and delayed retrieval, and (2) to assess whether sleep plays a role in the integration of two overlapping pairwise associations (i.e., open-loops).

Lau et al. (2010) found that when participants learnt overlapping A-B and A-C pairs, their ability to infer a relationship between B and C increased following a nap, relative to an awake condition. However, B-C inference can potentially occur via two means: (1) encoding generalisation; where prior to retrieval A-B and A-C associations are integrated into a generalised representation that potentially forms a direct association between B and C (Shohamy & Wagner, 2008; Zeithamova, Dominick, & Preston, 2012), or (2) retrieval generalisation, where the relationship between B and C is inferred ‘on-the-fly’ at the point of retrieval (Banino, Koster, Hassabis, & Kumaran, 2016; Kumaran & McClelland, 2012). Assessing retrieval accuracy and dependency for open-loops allows for differentiating between these two proposals; under the assumption that if A-B and A-C pairs are integrated prior to retrieval, dependency will be seen. Thus, if increases in B-C inference are seen as a function of sleep, with an associated increase in dependency, then this would support encoding-based generalisation. If B-C inference increases without any increase in dependency, then this would support retrieval-based generalisation. In the latter case, sleep might increase the associative strength of the directly encoded A-B and A-C pairs, and this might subsequently increase the probability of correct B-C inference at retrieval.

2.3.1. Methods

2.3.1.1. Participants

From previous published work (Horner et al., 2015; Horner & Burgess, 2013, 2014), with *n* = 177, an effect size *d* = .62 was calculated on the ability to detect a significant difference between the proportion of joint retrieval in the data and independent model (see 'Modelling retrieval dependency' section below). Using G*Power (Faul, Erdfelder, Buchner, & Lang, 2009), I conducted a power analysis with *d* = .62 and estimated that a sample size of *n* = 26 was required to detect a significant effect, if one is present, with a power of .85 and *a* = .05.

One hundred four participants (26 participants per condition, across 4 between-subject conditions) gave written informed consent to participate in Experiment 1. Participants were recruited from the University of York student population and took part in exchange for course credit or monetary compensation. Participants took part in one of four conditions: 26 participants in the open-loop, awake condition (23 female, mean age = 19.88 years, age range = 18-28 years), 26 in the open-loop, sleep condition (19 female, mean age = 19.68, age range = 18-23 years), 26 in the closed-loop, awake condition (23 female, mean age = 20.15 years, age range = 18-25), and 26 in the closed-loop, sleep condition (20 female, mean age = 20.65, age range = 18-28 years).

All experiments were approved by the Department of Psychology Ethics Committee, University of York, UK.

2.3.1.2. Materials

The stimuli consisted of 60 locations (e.g., *kitchen*), 60 famous people (e.g., *Barack Obama*) and 60 common objects (e.g., *hammer*). From these, 60 random location-person-object triplets were generated for each participant. Note that ‘triplet’ is used to refer to the three elements (location, person and object) that were assigned to the same associative structure (closed- or open-loop). Triplets were randomly assigned across the experimental conditions closed- vs open-loops. For closed-loops all three possible pairwise associations for a given triplet were encoded. For open-loops only two out of the three pairwise associations were encoded. Triplets were never presented all together at encoding or retrieval. Only specific pairwise associations were encoded and retrieved for each triplet, dependent on whether they were open- or closed-loops. Triplets were randomly assigned to the within-subject experimental conditions tested (i.e., tested at T1) vs not-tested (i.e., not tested at T1).

Note that the open-loop condition is equated to the closed-loop condition in the number of elements, but not in terms of the number of associations. It has previously been shown that a lack of dependency for open-loops is seen when three overlapping associations are encoded in an associative chain (e.g., *kitchen-hammer*, *kitchen-Barack Obama*, *Barack Obama-dog*), controlling for the number of associations (but not the number of elements) between open- and closed-loops (Horner & Burgess, 2014). Any differences in dependency between the two conditions in the current experiments are therefore unlikely to be driven by differences in the number of associations. Although the exposure to each pairwise association across open- and closed-loops is controlled for, two of the individual elements in the open-loop condition are only presented once, whereas all elements are presented twice in the closed-loop condition. Controlling for the exposure to each element would require repetition of pairwise associations in the open-loop condition. Here, I prefer to control for the number of exposures to each pairwise association (as this is what is being tested at retrieval), rather than the number of exposures to each individual element.

2.3.1.3. Procedure

The experiment consisted of a single encoding session and two retrieval sessions. Self-report ratings of alertness were collected before encoding and the second retrieval session using the

Stanford Sleepiness Scale (Hoddes, Zarcone, Smythe, Phillips, & Dement, 1973) (see Table 1 for mean sleepiness ratings). Ratings were not collected from four participants in the sleep, open-loop condition, and one participant in the sleep, closed-loop condition.

Table 1 Mean sleepiness ratings (and standard deviations) for retrieval sessions T1 and T2 for Experiment 1 for the sleep and awake condition. Note, the Stanford Sleepiness Scale ranges from 1 (feeling active, vital, alert, or wide awake) to 7 (no longer fighting sleep, sleep onset soon, having dream-like thoughts).

		Session	
Condition	**Loop**	**T1**	**T2**
Sleep	Open	2.23 (0.61)	2.23 (1.07)
-	Closed	3.00 (1.00)	2.64 (1.04)
Awake	Open	2.92 (1.23)	2.23 (1.14)
-	Closed	2.39 (0.80)	2.50 (1.36)

Session 1 (T1) took place between ~8-9 am for participants in the awake condition (open-loop: *M* = 8:31 am, range = 7:57-9:19 am) (closed-loop: *M* = 8:46 am, range = 8:04-9:31 am) and ~8-9 pm for participants in the sleep condition (open-loop: *M* = 8:52 pm, range = 8:05-9:34 pm; closed-loop: *M* = 8:48 pm, range = 7:48-9:34 pm). T1 consisted of a single encoding phase, and a retrieval phase (see details below). Participants in the awake conditions spent the remainder of the day normally and returned ~12 hours later for session 2 (T2) (open-loop: *M* = 11 hr, 50 min, range = 11 hr, 40 min-12 hr, 7 min; closed-loop: *M* = 11 hr, 49 min, range = 11 hr, 28 min-12 hr, 3 min) at ~8-9 pm. Participants in the post-encoding sleep condition returned to their own residence, slept overnight, and returned ~12 hours later (open-loop: *M* = 11 hr, 51 min, range = 11 hr, 23 min–12 hr, 16 min; closed-loop: *M* = 11 hr, 53 min, range = 11 hr, 44 min-12 hr, 13 min) at ~8-9 am. A 2x2 (Loop x Sleep) between-subjects ANOVA, with the factor Loop referring to whether participants encoded open- or closed-loops, and

Sleep referring to whether encoding was followed by sleep or wakefulness, revealed no significant difference in the duration of the interval between T1 and T2 across conditions, *F*s < 2.00, *p*s > .16.

Participants in the sleep condition completed a sleep diary prior to T1 and T2. Self-reported sleep durations were not collected from three participants in the open-loop condition and one participant in the closed-loop condition. Sleep quality ratings were not collected from three participants in the open- and closed-loop condition, respectively. There was no differences in self-reported duration, *t*(46) = .22, *p* = .83, *d* = .06, or quality, *t*(44) = .54, *p* = .59, *d* = .16, of sleep between T1 and T2 for participants in the open vs. closed-loop sleep conditions.

Encoding (T1). During encoding, participants were presented with specific pairwise associations for each of the 60 triplets. Participants learnt one pairwise association per trial. All pairwise associations were presented on a computer screen as words, with one item to the left and one to the right of fixation. The left/right assignment was randomly chosen on each trial. The words remained on screen for 6s. Participants were instructed to imagine, as vividly as possible, the items interacting in a meaningful way for the full 6s. Each word-pair presentation was preceded by a 500-ms fixation and followed by a 500-ms blank screen. In the open-loop condition, participants encoded, for each triplet, two (out of the three possible) pairwise associations; making a total of 120 encoding trials. For each triplet in the closed-loop condition, participants encoded all three pairwise associations; making a total of 180 encoding trials.

The encoding phase consisted of two or three blocks, for the open- and closed-loops respectively, of 60 trials with one pair from each triplet being presented during each block (participants were not made aware of this structure). A break of 20s would follow every 30 encoding trials. Within each block, the order of presentation was randomised. Each open-loop consisted of a common item (e.g., if the participants learnt location-person and then location-object, location would be the common item). 20 open-loops were pseudo-randomly assigned to each of the three possible common items (i.e., locations, people or objects). The presentation order for open-loops across the two blocks was: (1) person-location, location-object; (2) location-object, object-person; (3) object-person, person-location. Closed-loops were pseudo-randomly rotated in the same manner. The presentation order for the closed-loops across the three encoding blocks was: (1) person-location, location-object, object-person; (2) location-object, object-person, person-location; (3) object-person, person-location, location-object.

Retrieval (T1 and T2). During retrieval, participants performed a six-alternative forced-choice cued-recognition task. On a given trial, the cue and six possible targets were presented simultaneously on screen. The cue was presented in the middle of the screen with six possible targets; one target and five foils from the same category (e.g. if the target word was hammer, the five foils would be other randomly selected objects from other triplets), presented in two rows of three below the cue. Participants had 6s to respond with a key press and were instructed to be as accurate as possible in the time given. The location of the correct target item was randomly selected on each retrieval trial. Missing responses (*M* ± *SD* = 5.21 ± 7.01%) were treated as incorrect trials for the accuracy and dependency analyses. A 2x2 (Loop x Sleep) between-subjects ANOVA, where the dependent variable refers to the proportion of non-responses (collapsed across retrieval at T1 and T2), showed no significant effects, *F*s < 2.50, *p*s > .11.

For T1, 30 out of 60 triplets were tested. Each triplet was tested with one of the cue-target associations (e.g., cue: person, target: location) in both directions. For the open-loop condition, cue-target associations were presented across four blocks (with a single, randomly assigned, pairwise association from each triplet tested in each block), making a total of 120 trials. Only the directly encoded pairwise associations for open-loops were tested at T1 (i.e., no inference test was performed). For the closed-loop condition, the associations were presented across six blocks (i.e., three pairwise associations, tested in both directions, randomly assigned across blocks), making a total of 180 trials. A 20s break would follow every 30 trials. At T2, participants performed the same cued-recognition task as during T1 with all the triplets tested, making a total of 240 and 360 trials for the open- and closed-loop conditions, respectively.

For the open-loop condition, participants performed an additional inference test following the main cued-recognition task at T2. For example, if a participant had encoded the pairwise associations *Barack Obama-hammer* and *hammer-kitchen,* the non-encoded association *Barack Obama-kitchen* would be tested in both directions (i.e., cue *Obama,* retrieve *kitchen,* cue *kitchen,* retrieve *Obama*) during the inference task. For the inference task, the non-encoded associations for each open-loop were tested, in each direction, across two blocks, making a total of 120 trials. A 20s break would follow every 30 trials. Participants were not explicitly told that these were 'inference' trials and carried out the task in the same manner as for directly encoded pairs.

In the main analysis comparing retrieval at T1 vs T2, only retrieval trials at T2 for triplets that were not tested at T1 are included in order to control for possible testing effects. Further analyses that directly compare retrieval accuracy and dependency at T2 for triplets tested previously vs triplets not tested previously at T1 are also reported.

2.3.1.4. Modelling retrieval dependency

Six independent 2x2 contingency tables for the observed data and independent model were created for each participant in order to assess the dependency between the retrieval of two items (e.g., person, object) when cued by a common item (e.g., location) (A_BA_C), and between the retrieval of a common item (e.g., location) when cued by the other two items (e.g., person, object) (B_AC_A). Once constructed, the proportion of joint retrieval and joint non-retrieval in the data and independent model was calculated for each contingency table separately, by summing the leading diagonal cells and dividing by the total number of triplets (i.e., the proportion of triplets where two overlapping pairwise associations within a triplet are both retrieved either correctly or incorrectly). This measure was then averaged across the six contingency tables to provide a single measure of the proportion of joint retrieval and non-retrieval for the data and independent model separately. I refer to this measure as the 'proportion of joint retrieval' but note that it includes both the proportion of joint retrieval and joint non-retrieval.

Table 2 Contingency table for the independent model for correct and incorrect retrieval, over *N* triplets (*i* = 1 to *N*), for elements B and C when cued by A.

Retrieval of element C	**Retrieval of element B**	
	Correct (P_{AB})	Incorrect (1-P_{AB})
Correct (P_{AC})	$\sum_{i=1}^{N} P_{AB_i} P_{AC_i}$	$\sum_{i=1}^{N} P_{AC_i}(1-P_{AB_i})$
Incorrect (1-P_{AC})	$\sum_{i=1}^{N} P_{AB_i}(1-P_{AC_i})$	$\sum_{i=1}^{N} (1-P_{AB_i})(1-P_{AC_i})$

The independent model assumes that pairwise associations for a given triplet are retrieved independently of one another – i.e., if a participant retrieves one pairwise association from a triplet (un)successfully this does not predict the participants ability to (un)successfully retrieve another pairwise association from the same triplet. As such, the independent model serves as a lower bound that can be compared with the proportion of joint retrieval in the data. Note that the proportion of joint retrieval measure scales with accuracy, and as such only comparisons between the data and independent model (i.e., the 'dependency' measure reported below) are meaningful.

The 2x2 contingency tables for the data show the number of triplets that fall within the four cells (i.e., for the $A_B A_C$ analysis, both A_B and A_C correct; A_B incorrect and A_C correct; A_B correct and A_C incorrect; and both A_B and A_C incorrect, where A_B = cue with location (A) and retrieve person (B) and similarly for A_C, where C stands for object). The table for the independent model (Table 2) shows the predicted proportion of triplets that fall in the four cells, given a participant's overall level of accuracy, if the retrieval of within-event associations is assumed to be independent.

For a given participant, the proportion of correct retrievals of, for instance, item B when cued by A is denoted by P_{AB} (i.e., the mean performance for B when cued by A across all events). For the independent model, when cued by A, the probability of (1) correctly retrieving both B and C (across all triplets) is equal to P_{AB} P_{AC}; (2) correctly retrieving B but not C is equal to P_{AB} (1 - P_{AC}); (3) correctly retrieving C but not B is equal to (1 - P_{AB}) P_{AC}; and (4) incorrectly retrieving both B and C is equal to (1 - P_{AB}) (1 - P_{AC}).

2.3.1.5. Statistical analyses

For the main analysis of retrieval accuracy, I report a 2x2x2 (Session x Loop x Sleep) mixed ANOVA with the within-subject factor Session referring to T1 (immediate) vs T2 (12-hour delay), the between-subject factor Loop referring to whether participants encoded open- vs closed-loops, and the between-subject factor Sleep referring to whether T1 was followed by sleep vs wakefulness. The main analysis reports memory performance for triplets at T2 that were not previously tested at T1. A 2x2x2 (Tested x Loop x Sleep) mixed ANOVA is also reported for retrieval accuracy at T2, with the within-subject factor Tested referring to whether triplets had previously been tested at T1 or not.

For the main dependency analysis, I report a 2x2x2 (Session x Loop x Sleep) mixed ANOVA, where the dependent variable refers to the difference between the proportion of joint retrieval in the data and independent model (i.e., ‘dependency’). A 2x2x2 (Tested x Loop x Sleep) mixed ANOVA is also reported for retrieval dependency at T2, with the within-subject factor Tested again referring to whether triplets had previously been tested at T1 or not. I also report *t*-tests comparing the proportion of joint retrieval in the data with their respective independent models.

Alpha was set to .05 (two-tailed) for all statistical tests. For each ANOVA, a partial eta-squared effect size (η_p^2) is reported. For *t*-tests, I report a Cohen's *d* as the mean difference between the conditions divided by the pooled standard deviation across conditions (Lakens, 2013) as an estimate of the between-subjects effect size (regardless of whether the effect is within- or between-subjects). All statistical analyses were conducted using JASP (jasp-stats.org; JASP Team, 2018).

2.3.2. Results

2.3.2.1. Retrieval accuracy

Mean proportion correct (and standard deviations) across Session, Loop and Sleep are presented in Table 3, and mean proportion correct across Loop and Session (collapsed across Sleep) is presented in Figure 3.

A 2x2x2 (Session x Loop x Sleep) ANOVA revealed a significant effect of Session, with retrieval accuracy decreasing from T1 to T2, $F(1,100) = 352.02$, $p < .001$, $\eta_p^2 = .78$. The ANOVA also revealed a significant interaction between Session and Sleep, $F(1,100) = 59.06$, $p < .001$, $\eta_p^2 = .37$, with significantly more forgetting between T1 and T2 in the awake relative to sleep condition. No significant main effects were seen for Loop, $F(1,100) = .16$, $p = .69$, $\eta_p^2 < .01$, and Sleep, $F(1,100) = 3.88$, $p = .05$, $\eta_p^2 = .04$ (but note the borderline *p*-value for the main effect of Sleep). A further interaction between Session and Loop was also seen, $F(1,100) = 42.46$, $p < .001$, $\eta_p^2 = .30$, suggesting possible greater forgetting for open- than closed-loops. This interaction appeared to occur regardless of Sleep, given there was no interaction between Session, Loop, and Sleep, $F(1,100) = .15$, $p = .70$, $\eta_p^2 < .01$.

Table 3 Mean proportion correct (and standard deviations) at retrieval sessions T1 and T2 for Experiment 1-4. For T2, only trials where participants retrieved cue-target associations not previously tested at T1 are included. n/a = not applicable.

			Session	
	Condition	**Loop**	**T1**	**T2**
Experiment 1	Sleep	Open	.73 (.16)	.61 (.16)
	-	Closed	.72 (.25)	.69 (.27)
	Awake	Open	.73 (.14)	.51 (.14)
	-	Closed	.68 (.21)	.54 (.23)
Experiment 2	n/a	Open	.71 (.17)	.41 (.14)
	n/a	Closed	.74 (.17)	.46 (.20)
Experiment 3	n/a	Open	n/a	.35 (.15)
	n/a	Closed	n/a	.51 (.24)
Experiment 4	n/a	Open	.69 (.18)	.37 (.14)
	n/a	Closed	.73 (.18)	.50 (.19)

A 2x2x2 (Tested x Loop x Sleep) ANOVA revealed a main effect of Tested, $F(1,100) = 408.87$, $p < .001$, $\eta_p^2 = .80$, with greater retrieval accuracy for triplets previously tested at T1. A interaction between Tested and Sleep was also seen, $F(1,100) = 15.16$, $p < .001$, $\eta_p^2 = .13$, with a greater effect of sleep for triplets not previously tested, suggesting that prior testing decreased the effect of sleep. An interaction between Tested and Loop was also seen, $F(1,100) = 13.69$, $p < .001$, $\eta_p^2 = .12$, with a greater testing effect for open- relative to closed-loops. The ANOVA also revealed a main effect of Sleep, $F(1,100) = 6.54$, $p = .01$, $\eta_p^2 = .06$, consistent with the interaction between Session and Sleep seen in the main analysis above. No further significant effects or interactions were seen, *F*s < 2.58, *p*s >.11. Prior testing therefore increased performance at T2 relative to associations not previously tested at T1, and this effect was modulated by Sleep and Loop.

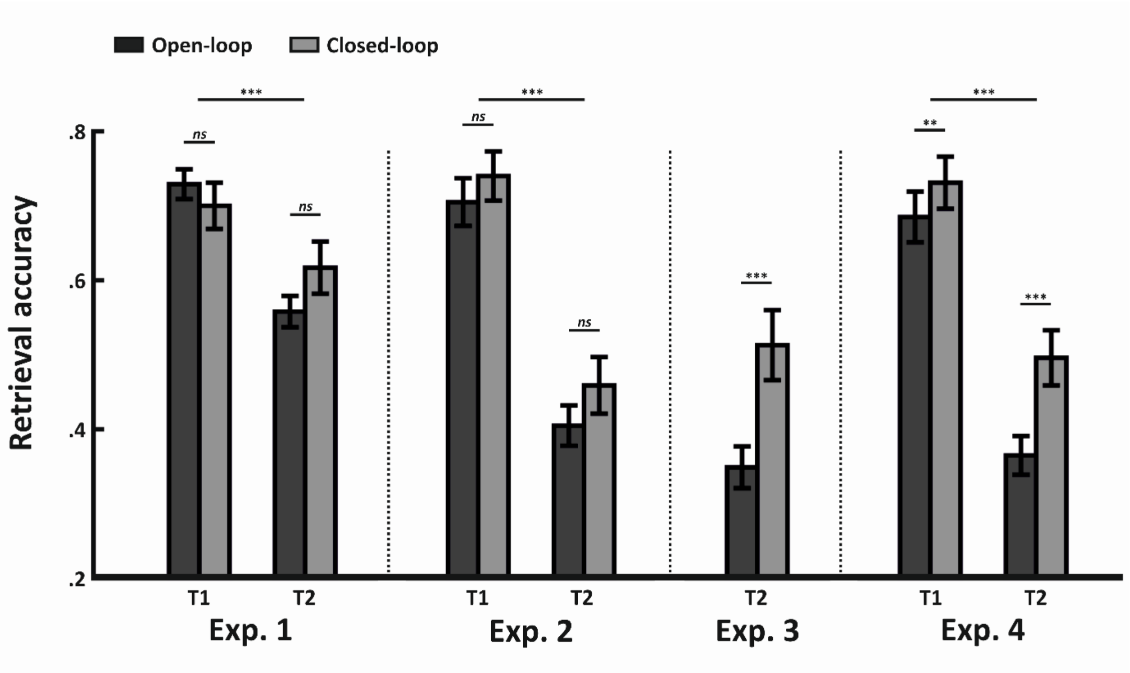

Figure 3 Mean proportion correct for open- and closed-loops at retrieval sessions T1 and T2 for Experiments 1-4. For T2, only trials where participants retrieved triplets not previously tested at T1 are included. Error bars represent ± 1 standard error. Exp = Experiment. *** $p < .001$, ** $p < .01$.

2.3.2.2. Retrieval dependency

Mean proportion of joint retrieval (and standard deviations) for the data and independent model for open- and closed-loops (collapsed across Sleep), are presented in Table 4. Figure 4 shows dependency across Session and Loop (collapsed across Sleep).

Consistent with previous research (Horner & Burgess, 2014), no evidence of dependency was seen for open-loops at T1, $t(51) = 1.63$, $p = .11$, $d = .09$, but dependency was seen for closed-loops at T1, $t(51) = 6.03$, $p < .001$, $d = .20$. Critically, closed-loops still showed dependency at T2, $t(51) = 5.31$, $p < .001$, $d = .23$, while significant anti-dependency was seen for open-loops, $t(51) = 2.48$, $p = .02$, $d = .26$ (but note that this effect does not survive correction for multiple comparisons).

Table 4 Mean proportion of joint retrieval (and standard deviations) for the data and independent model for retrieval sessions T1 and T2 for Experiment 1-4. For Experiment 1, the proportion of joint retrieval is collapsed across the sleep and awake conditions. For T2, only trials where participants retrieved cue-target associations not previously tested at T1 are included. n/a = not applicable.

		Session			
		T1		**T2**	
	Loop	**Data**	**Independent**	**Data**	**Independent**
Experiment 1	Open	.62 (.14)	.64 (.12)	.52 (.10)	.55 (.06)
	Closed	.71 (.17)	.68 (.18)	.69 (.15)	.66 (.17)
Experiment 2	Open	.60 (.16)	.61 (.15)	.56 (.08)	.55 (.08)
	Closed	.73 (.14)	.67 (.16)	.62 (.08)	.58 (.09)
Experiment 3	Open	n/a	n/a	.57 (.10)	.58 (.08)
	Closed	n/a	n/a	.69 (.09)	.61 (.11)
Experiment 4	Open	.57 (.16)	.59 (.12)	.56 (.10)	.56 (.08)
	Closed	.70 (.15)	.66 (.16)	.60 (.09)	.57 (.07)

Anti-dependency in the open-loops suggests that associations interfere with each other. It is possible that anti-dependency emerges during the retention period between T1 and T2, as anti-dependency is not observed at T1. Although no significant anti-dependency was seen at T1, it is possible that the associations already interfere with each other either at immediate retrieval, or at the point of encoding, given that a lower proportion of joint retrieval in the data, relative to the independent model, is observed immediately after encoding. Consistent with this, the main analysis showed that dependency did not change significantly between T1 and T2 across the closed- and open-loop conditions. Note that this anti-dependency effect is not replicated in Experiments 2-4.

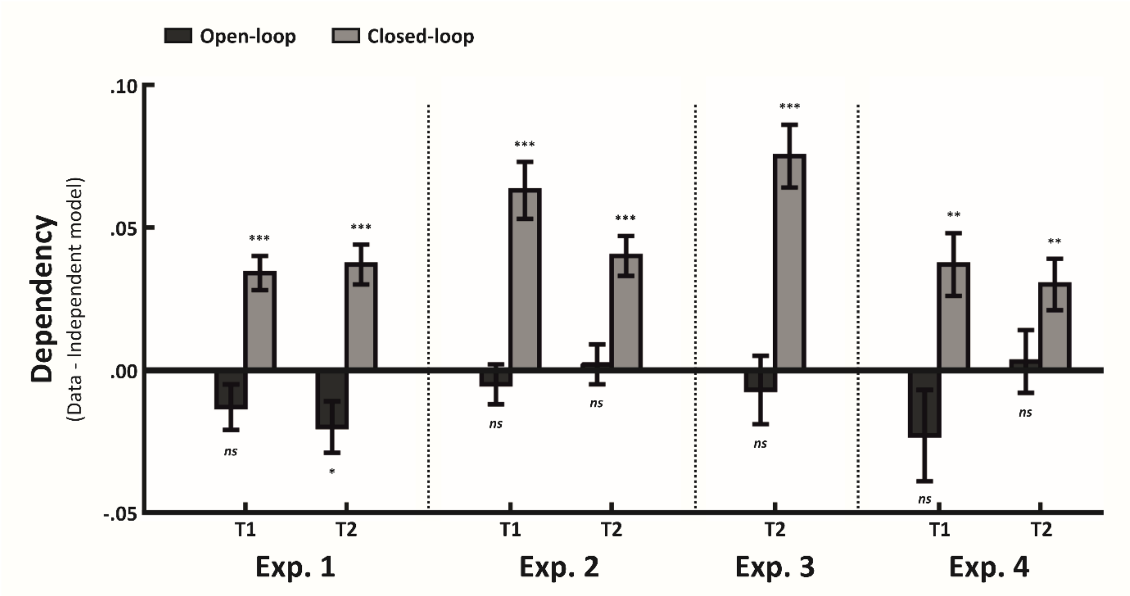

Figure 4 Dependency for open- and closed-loops at retrieval sessions T1 and T2 for Experiments 1-4. For T2, only trials where participants retrieved triplets previously not tested at T1 are included. Error bars represent ± 1 standard error. Exp = Experiment. *ns* = not significant. *** $p < .001$, ** $p < .01$, * $p < .05$ (uncorrected).

A 2x2x2 (Session x Loop x Sleep) ANOVA on the dependency revealed a significant main effect of Loop, $F(1,100) = 37.02$, $p < .001$, $\eta_p^2 = .27$, confirming significantly greater dependency for closed- than open-loops. There was no evidence for changes in dependency across Session, $F(1,100) = .25$, $p = .62$, $\eta_p^2 < .01$, nor did Session interact with Sleep, $F(1,100) = .16$, $p = .70$, $\eta_p^2 < .01$, or Loop, $F(1,100) = .94$, $p = .34$, $\eta_p^2 < .01$. Indeed, no other significant effects or interactions were seen, *F*s < 1.17, *p*s > .28, beyond the main effect of Loop. A 2x2x2 (Tested x Loop x Sleep) ANOVA also revealed a significant main effect of Loop, $F(1,100) = 44.29$, $p < .001$, $\eta_p^2 = .31$. No other significant effects or interactions were seen, *F*s < 1.89, *p*s > .17. Accordingly, despite large variation in retrieval performance at T2, relative to T1, as a function of sleep (and prior testing), dependency in the closed-loop and open-loop condition remained consistent between T1 and T2.

2.3.2.3. Control analyses

Retrieval dependency for closed-loops including only two pairwise associations. It has previously been shown that closed-loops show significantly greater dependency than open-loops, even when the number of learnt elements or associations between closed- and open-loops is equated (Horner & Burgess, 2014; Horner et al., 2015). For consistency, I repeated the main dependency analysis for closed-loops, but included only two out of the three pairwise associations (e.g., *Barack Obama-kitchen, kitchen-hammer,* but not *Barack Obama-hammer*). As such, the number of pairwise associations included in the dependency analysis for closed-loops is equivalent to that included for open-loops reported in the main analysis. Consistent with previous published work (Horner & Burgess, 2014), despite including only two out of the three learnt associations, closed-loops show significant dependency at both T1, $t(51) = 4.78$, $p < .001$, $d = .16$ and T2, $t(51) = 4.94$, $p < .001$, $d = .21$. Similarly, a 2x2x2 (Session x Loop x Sleep) ANOVA replicated the significant effect of Loop, $F(1,100) = 31.29$, $p < .001$, $\eta_p^2 = .24$, reported in the main analysis. No other significant main effects or interactions were observed, *F*s < 2.26, *p*s > .13. A 2x2x2 (Tested x Loop x Sleep) ANOVA also replicated the main effect of Loop, $F(1,100) = 36.33$, $p < .001$, $\eta_p^2 = .27$, in addition to revealing a significant interaction between Tested and Sleep, $F(1,100) = 4.72$, $p = .03$, $\eta_p^2 = .05$, with a greater difference in dependency for previously tested vs not tested triplets in the awake, relative to the sleep condition.

Retrieval dependency across element- and analysis-type. In previous work, Horner and Burgess (2014) showed that the proportion of joint retrieval in the observed data does not vary significantly across element-type (i.e., people, locations, and objects), and analysis type (A_BA_C, where the element A refers to the common cue element-type; and B_AC_A, where the element A refers to the common retrieved element-type). In order to assess for differences in dependency that might be due to variations in the memorability of element-types across triplets, I repeated the main analysis across individual element-types for the two separate analysis types (i.e., $A_B A_C$ and $B_A B_C$).

A 2x2x3 (Session x Loop x Cue-type) ANOVA on dependency, collapsed across retrieval-type (i.e., the different element-types in the $B_A B_C$ analysis) and the sleep and awake condition, replicated the main effect of Loop, $F(1,102) = 20.7$, $p < .001$, $\eta_p^2 = .17$. No other main effects or interactions were observed, *F*s < 1.87, *p*s > .15. Note that cue-type refers to the different element-types (i.e., people, locations, and objects) in the $A_B A_C$ analysis.

Similarly, a 2x2x3 (Session x Loop x Retrieval-type) ANOVA on dependency, collapsed across cue-type (i.e., the different element-types in the $A_B A_C$ analysis) and the sleep and awake condition, replicated the main effect of Loop, $F(1,102) = 36.86$, $p < .001$, $\eta_p^2 = .27$. The ANOVA did not reveal any other significant main effects or interactions, *F*s < 2.17, *p*s > .11, although note a borderline interaction between Session and Loop, $F(1,102) = 3.42$, $p < .07$, $\eta_p^2 = .32$. Note that retrieval-type refers to the different element-types (i.e., people, locations, and objects) in the $B_A C_A$ analysis.

2.3.2.4. Mnemonic integration during sleep

I was also interested in assessing the possible role that sleep may play in integrating overlapping information. For this analysis, I focused solely on the open-loops as these are equivalent to the A-B, A-C structures encoded in Lau et al. (2010). The main analysis found an overall effect of sleep on accuracy (that did not interact with Loop), but no effect on dependency. Here, given the specific interest in whether sleep modulates mnemonic inference, I only report analyses for the open-loop condition (as no inference is possible for closed-loops).

For open-loop retrieval accuracy, a 2x2 (Session x Sleep) ANOVA revealed a significant interaction between Session and Sleep, $F(1,50) = 25.30$, $p < .001$, $\eta_p^2 = .34$, consistent with the main analysis showing that sleep decreases forgetting for directly encoded pairs.

Participants' ability to infer non-encoded B-C pairs at T2 (see Table 5) was also assessed. Note that I did not assess B-C inference at T1 as this may have increased participants' awareness of the relationship between the overlapping pairs, possibly leading to increases in dependency.

Table 5 Mean proportion correct (and standard deviations) for encoded cue-target associations at T1 and T2, and non-encoded pairs at T2 for the open-loop condition in Experiment 1. For T2, only trials where participants retrieved cue-target associations not previously tested at T1 are included.

	Encoded		**Non-encoded**
Condition	**T1**	**T2**	**T2**
Sleep	.73 (.16)	.61 (.16)	.53 (.22)
Awake	.73 (.14)	.51 (.14)	.42 (.29)

One participant in the sleep condition was excluded from this analysis due to a failure to respond during the inference task (missing responses > 80%). Accordingly, 51 participants (26 in the awake condition, and 25 in the sleep condition) were included in the analysis.

Table 6 Mean proportion of joint retrieval (and standard deviations) for the data and independent model for retrieval sessions T1 and T2 for the open-loop condition in Experiment 1. For T2, only trials where participants retrieved cue-target associations not previously tested at T1 are included.

	Session			
	T1		**T2**	
Condition	**Data**	**Independent**	**Data**	**Independent**
Sleep	.62 (.14)	.64 (.11)	.54 (.10)	.56 (.07)
Awake	.62 (.15)	.63 (.14)	.51 (.09)	.53 (.05)

Consistent with Lau et al. (2010), greater B-C inference performance was seen in the sleep, relative to awake condition, $t(49) = 2.03$, $p = .048$, $d = .57$ (though note the borderline p-value). Importantly, a 2x2 (Session x Sleep) ANOVA failed to show any evidence for a change in dependency for open-loops between T1 and T2, $F(1,50) = .72$, $p = .40$, $\eta_p^2 = .01$ (see Table 6), nor did Session interact with Sleep, $F(1,50) = .16$, $p = .70$, $\eta_p^2 < .01$. Thus, evidence for increases for inference performance was seen, but there was no evidence for an increase in dependency as a function of sleep.

2.3.3. Discussion

Experiment 1 modulated retrieval accuracy by manipulating (1) the time between encoding and retrieval, (2) whether participants slept between encoding and retrieval, and (3) whether pairwise associations were previously tested or not. Evidence for effects of all three manipulations on retrieval accuracy was seen, such that large variations in the amount of forgetting were observed across conditions. Despite this, I saw no evidence for changes in dependency for either open- or closed-loops. No dependency (or anti-dependency) was seen for open-loops, and dependency was consistently seen for closed-loops. Experiment 1 therefore provides evidence that dependency does not change over time – i.e., closed-loops retain their dependency while open-loops do not show dependency.

No evidence for mnemonic integration during sleep (as measured by retrieval dependency) was seen, suggesting that the role sleep plays in increasing inference is unlikely to be driven by encoding generalisation during sleep. It is perhaps more likely that inference is driven by ‘on-the-fly’ processes at the point of retrieval; the probability of which is increased due to less forgetting for directly encoded pairs in the sleep than awake condition. Although I saw no evidence for increases in dependency for open-loops following sleep, participants were able to make the correct inferences at a level well above chance. The task instructions were ambiguous in relation to the inference task - i.e., participants were presented with inference trials as if they were retrieval trials. Further work is needed to clarify whether participants were making correct inferences based on false memories for non-encoded pairs, or whether they were making informed inferences. However, the lack of dependency following sleep suggests that this inference process is likely to be occurring at the point of retrieval.

2.4. Experiment 2

Despite evidence for dependency in the closed-loop condition after 12 hours in Experiment 1, it is possible that the amount of forgetting in Experiment 1 was not sufficient to produce fragmentation, and in turn decreases in dependency. In Experiments 2-4, I therefore tested participants after a week, rather than a 12-hour delay. This extended interval between T1 and T2 produced greater (numerical) amounts of forgetting relative to Experiment 1, creating a sterner test for the hypothesis that closed-loops are forgotten in an all-or-none manner.

2.4.1. Methods

Experiment 2 was identical to Experiment 1 with the following exceptions. Experiment 2 equated to a 2x2 design with the factors Session and Loop. No factor of Sleep was included, given the interval between encoding and retrieval was 1 week.

2.4.1.1. Participants

Fifty-two participants gave written informed consent to participate in Experiment 2. Participants were recruited from the University of York student population. Participants took part in exchange for course credit or monetary compensation, and were allocated to one of two conditions: 26 participants in the open-loop condition (23 female, mean age = 19.46 years, age range = 18-23) and 26 participants in the closed-loop condition (23 female, mean age = 20.00 years, age range = 18-26).

2.4.1.2. Procedure

In order to increase the amount of forgetting relative to Experiment 1, the two sessions were separated by one week. All sessions took place in the afternoon. Encoding and T1 took place between ~12-5 pm (open-loop: M = 2:31 pm, range = 12:00-4:46 pm; closed-loop: M = 2:15 pm, range = 11:58 am-4:47 pm). T2 took place 1 week later between ~12-5 pm (open-loop: M = 2:29 pm, range = 11:51 am-4:41 pm; closed-loop: M = 2:08 pm, range = 11:58 am-4:41 pm). Missing responses during retrieval ($M \pm SD = 4.42 \pm 4.88\%$) were again treated as incorrect trials. There was no difference in the proportion of missed responses (collapsed across Session) between open- and closed-loops, $t(50) = .89$, $p = .38$, $d = .25$.

2.4.2. Results

2.4.2.1. Retrieval accuracy

Mean proportion correct across conditions are shown in Table 3 and Figure 3. Retrieval accuracy was .72 at T1 and .43 at T2. This is compared to retrieval accuracy of .72 at T1 and .59 at T2 in Experiment 1. Thus, increasing the interval between T1 and T2 to one week led to numerically greater forgetting relative to a 12-hour interval. A 2x2 (Session x Loop) ANOVA revealed a main effect of Session, $F(1,50) = 318.83$, $p < .001$, $\eta_p^2 = .86$, confirming a significant decrease in performance at T2 relative to T1. No further effects or interactions were seen, Fs < 1.01, ps > .32. A 2x2 (Tested x Loop) ANOVA on retrieval accuracy at T2 revealed a main effect of Tested, $F(1,50) = 187.48$, $p < .001$, $\eta_p^2 = .79$, with greater accuracy for previously tested than not previously tested associations. No other significant main effects or interactions were observed, Fs < 1.87, ps > .17.

2.4.2.2. Retrieval dependency

Mean proportion of joint retrieval (and standard deviations) for the data and independent model across conditions are presented in Table 4. Figure 4 shows dependency across Session and Loop. As in Experiment 1, no evidence for dependency was observed for open-loops at T1, $t(25) = .67$, $p = .51$, $d = .03$, or T2, $t(25) = .28$, $p = .78$, $d = .03$, but significant evidence for dependency was seen for closed-loops at both T1, $t(25) = 5.90$, $p < .001$, $d = .42$, and T2, $t(25) = 5.31$, $p < .001$, $d = .51$. A 2x2 (Session x Loop) ANOVA on dependency revealed a significant effect of Loop, $F(1,50) = 39.96$, $p < .001$, $\eta_p^2 = .44$, confirming that dependency was significantly greater in the closed- relative to open-loop condition. No interaction between Session and Loop was seen, $F(1,50) = 3.34$, $p = .07$, $\eta_p^2 = .06$ (but note the marginal *p*-value).

In order to interrogate this marginal interaction further, I performed a *t*-test between dependency at T1 and T2 separately for each loop type. Consistent with the main analyses, no evidence for a change in dependency was seen for either closed-loops, $t(25) = 1.72$, $p = .10$, $d = .49$, or open-loops, $t(25) = .733$, $p = .47$, $d = .20$, between T1 and T2. Critically, as shown above, dependency was still significant in the closed-loop condition at T2. Thus, as in Experiment 1, despite high levels of forgetting between T1 and T2, no evidence for a decrease in dependency between T1 and T2 was seen for closed-loops.

A 2x2 (Tested x Loop) ANOVA on dependency at T2 revealed a significant main effect of Loop, $F(1,50) = 78.37$, $p < .001$, $\eta_p^2 = .61$, in addition to a significant interaction between Tested and Loop, $F(1,50) = 14.44$, $p < .001$, $\eta_p^2 = .22$, with a greater difference in dependency between closed- and open-loops previously tested at T1, relative to those not tested at T1.

2.4.3. Discussion

Experiment 2 produced greater numerical amounts of forgetting following an interval of one week between encoding and delayed retrieval, relative to Experiment 1. Despite this increase in forgetting (as measured by retrieval accuracy), I again saw no decrease in dependency for closed-loops, nor any increase in dependency for open-loops. Experiment 1 and 2 showed that forgetting can be affected by several post-encoding factors, such as the interval between encoding and retrieval, post-encoding sleep, and prior testing. However, across these factors, there is no evidence for decreases in dependency for closed-loops.

2.5. Experiment 3

Retrieval accuracy is typically greater for closed- than open-loops when both conditions are learnt within-subject (i.e., each participant learns both closed- and open-loops; Horner et al., 2015). In Experiment 1 and 2, there was no evidence to suggest that retrieval accuracy was higher for closed- than open-loops. In Experiment 3, I aimed to assess whether the lack of a difference in retrieval accuracy between closed- and open-loops in Experiments 1 and 2 was a function of the between-subjects design. This is theoretically important because if closed-loops are associated with higher accuracy relative to open-loops in a within-subject, but not between-subject manipulation, it might suggest a possible competitive mechanism between memory representations. In Experiment 3, participants learnt both closed- and open-loops at T1 and were tested in a single session (T2) after a week. Note that no immediate retrieval test was performed in order to ensure that the overall number of triplets per condition was consistent across Experiments 2 and 3 (30 per condition). Experiment 3 also provided a further opportunity to replicate the pattern of dependency seen for closed- and open-loops over the course of a week.

2.5.1. Methods

Experiment 3 was identical to Experiment 2 with the following exceptions.

2.5.1.1. Participants

Twenty-six participants (22 female, mean age = 19.35, age range = 18-23) gave written informed consent to participate in Experiment 3. Participants were recruited from the

University of York student population and took part in exchange for course credit or monetary compensation.

2.5.1.2. Materials

60 random location-person-objects triplets were generated for each participant. 30 triplets were randomly assigned to the within-subject open- and closed-loop conditions.

2.5.1.3. Procedure

The two sessions were separated by one week. All sessions took place in the afternoon. Encoding took place between ~12-5 pm (*M* = 2:24 pm, range = 11:57 am–4:41 pm). T2 took place one week later between ~12-5 pm (*M* = 2:23 pm, range = 11:56 am–4:39 pm).

Encoding. Participants were presented with specific pairwise associations for each of the 60 triplets. For 30 out of the 60 triplets, participants encoded all three possible pairwise associations; forming closed-loops. For the other 30 triplets, participants encoded two out of the three possible pairwise associations; forming open-loops. The encoding phase consisted of three blocks of 30, 60, and 60 trials, making a total of 150 encoding trials. During the first block, only pairwise associations for closed-loops were presented. This ensured that the duration between the encoding of the last pairwise association and T2 was consistent across closed- and open-loops. In blocks 2 and 3, open- and closed-loop associations were presented randomly in an intermixed manner.

Retrieval (T2). No immediate retrieval test followed encoding. This was done in order to maintain consistency in the number of closed- and open-loops tested at T2 in Experiments 2 and 3. At T2, all 60 triplets were tested. Cue-target associations were presented across six blocks, making a total of 300 trials. Note that none of these cue-target associations had been tested previously at T1. As at encoding, open- and closed-loops were presented randomly within each block. As open-loops were formed of only two out of the three possible pairwise associations, the four possible cue-target associations per open-loop were randomly distributed across four out of the six blocks. Note that this necessitates that the number of trials per block can vary between participants. Missing responses (*M* ± *SD* = 4.00 ± 7.04%) were treated as incorrect trials. There was no difference between open- and closed-loops in the proportion of missing responses, *t*(25) = 1.97, *p* = .06, *d* = .16 (though note the borderline p-value indicating possibly greater proportions of missing responses in the open- than closed-loop condition).

2.5.1.4. Statistical analyses

For retrieval accuracy, I report a paired samples *t*-test comparing performance for closed- vs open-loops. For retrieval dependency, I report a paired samples *t*-test comparing the proportion of joint retrieval for the data and independent model for closed- vs open-loops.

2.5.2. Results

2.5.2.1. Retrieval accuracy

Mean proportion correct for open- and closed-loops are shown in Table 3 and Figure 3. Retrieval accuracy at T2 was .43 (averaged across open- and closed-loops). This is comparable to .43 in Experiment 2. Accordingly, Experiment 3 produced numerically similar performance at T2 as Experiment 2. However, in contrast to Experiment 2, accuracy for closed-loops (.51) was greater than for open-loops (.35) following a one week delay, *t(*25) = 6.31, *p* < .001, *d* = .84.

2.5.2.2. Retrieval dependency

Consistent with Experiments 1-2, dependency was greater for closed- than open-loops, *t*(25) = 5.35, *p* < .001, *d* = 1.33, with closed-loops again showing a significantly greater proportion of joint retrieval in the data than in the independent model, *t*(25) = 6.40, *p* < .001, *d* = .73, and open-loops showing no evidence for dependency, *t*(25) = .56, *p* = .58, *d* = .08 (see Table 3 and Figure 4).

2.5.3. Discussion

Experiment 3 replicated the pattern of dependency seen in Experiments 1-2. Presuming a significant amount of forgetting occurred in Experiment 3, as seen in Experiments 1-2, I again observed that dependency for closed-loops is resilient to forgetting. Interestingly, a significant difference in retrieval accuracy between closed- and open-loops was seen; a pattern not seen in Experiment 2. The critical difference between Experiments 2 and 3 is that the Loop manipulation was a between-subject factor in Experiment 2, but a within-subject factor in Experiment 3.

However, Experiment 3 did not include an immediate test (in contrast to Experiment 2). It is unclear whether this difference between closed- vs open-loops in a within-subject design would also be present at T1. Therefore, in Experiment 4, the Loop manipulation was again a within-subject factor, however Experiment 4 also included an immediate, as well as delayed, retrieval phase. Note that this decreased the number of triplets per condition from 30 to 15.

2.6. Experiment 4

Experiment 4 included a retrieval test at both T1 and T2. This allowed for assessing if the difference in retrieval accuracy for closed- vs open-loops at T2 in Experiment 3 was also present at T1. It also presented an opportunity to replicate the retrieval accuracy difference at T2 seen in Experiment 3. Experiment 4 also offered an opportunity to replicate the pattern of dependency for closed-loops seen in Experiments 1-3.

2.6.1. Methods

Experiment 4 was identical to Experiment 3 with the following exceptions.

2.6.1.1. Participants

Twenty-seven participants gave written informed consent to participate in Experiment 4. Participants were recruited from the University of York student population and took part in exchange for course credit or monetary compensation. One participant was excluded due to a failure to respond at T2 (missing responses > 50%). Accordingly, 26 participants (25 female, mean age = 19.27, age range = 18-23) were included in the analyses.

2.6.1.2. Procedure

Encoding took place between ~12-5 pm (*M* = 2:26 pm, range = 12:02-4:56 pm). T2 took place one week later between ~12-5 pm (*M* = 2:26 pm, range = 11:47 am-16:48 pm). For T1, 30 out of 60 triplets were tested. Triplets were randomly assigned to the within-subject condition tested (i.e., tested at T1) vs not-tested (i.e., not tested at T1). As such, I could assess retrieval accuracy and dependency for 15 open-loops and 15 closed-loops immediately after encoding. Cue-target associations were presented across 6 blocks, making a total of 150 trials. A break of 20s followed every 25 trials. At T2, all 60 triplets were tested. Cue-target associations were presented across 6 blocks, making a total of 300 trials. A 20s break followed every 30 trials. The four possible cue-target associations per open-loop were randomly distributed across the six blocks. Missing responses (*M* ± *SD* = 5.73 ± 6.54%) were treated as incorrect trials. There was no difference between open- and closed-loops (collapsed across Session) in the proportion of missing responses, $t(25) = 1.60$, $p = .12$, $d = .20$.

2.6.1.3. Statistical analyses

For the main analysis of retrieval accuracy, a 2x2 (Session x Loop) within-subject ANOVA is reported. I also report a 2x2 (Tested x Loop) within-subject ANOVA for retrieval accuracy at T2, where the within-subject factor Tested refers to whether the triplets had previously been tested at T1 or not. For the dependency analysis, I report a 2x2 (Session x Loop) within-subject ANOVA and a 2x2 (Tested x Loop) within-subject ANOVA, where the dependent variable refers to the difference between the proportion of joint retrieval in the data and independent model.

2.6.2. Results

2.6.2.1. Retrieval accuracy

Mean proportion correct across conditions are shown in Table 3 and Figure 3. Retrieval accuracy was .71 at T1 and .43 at T2. This is consistent with performance seen in Experiments 2 and 3. A 2x2 (Session x Loop) ANOVA revealed a main effect of Session, $F(1,25) = 182.14$, $p < .001$, $\eta_p^2 = .88$, in addition to a significant main effect of Loop, $F(1,25) = 27.61$, $p < .001$, $\eta_p^2 = .53$, with greater accuracy for closed- relative to open-loops at both T1, $t(25) = 2.86$, $p < .01$, d = .26, and T2, $t(25) = 5.12$, $p < .001$, $d = .79$. A significant interaction between Session and Loop was also observed, $F(1,25) = 10.40$, p < .01, $\eta_p^2 = .29$, with the difference between closed- and open-loops increasing from T1 to T2. Thus, closed-loops show both higher retrieval accuracy (regardless of the retention interval) and higher levels of retention relative to open-loops.

A 2x2 (Tested x Loop) ANOVA on retrieval accuracy at T2 revealed a main effect of Tested, $F(1,25) = 142.12$, $p < .001$, $\eta_p^2 = .85$, with greater accuracy for previously tested relative to closed- and open-loops not tested previously. The ANOVA also revealed a main effect of Loop, $F(1,25) = 50.03$, $p < .001$, $\eta_p^2 = .67$, with overall greater retrieval accuracy for closed-loops than open-loops.

2.6.2.2. Retrieval dependency

Mean proportion of joint retrieval (and standard deviations) for the data and independent model across conditions are presented in Table 4. Dependency across Session and Loop is shown in Figure 4.

Consistent with Experiments 1-3, no evidence for dependency was seen for open-loops at T1, $t(25) = 1.39$, $p = .18$, $d = .16$, and T2, $t(25) = .24$, $p = .81$, $d = .03$. Significant dependency was seen for closed-loops at both T1, $t(25) = 3.29$, $p < .01$, $d = .25$, and T2, $t(25) = 3.21$, $p < .01$, $d = .38$.

A 2x2 (Session x Loop) within-subject ANOVA on dependency revealed a significant main effect of Loop, $F(1,50) = 20.20$, $p < .001$, $\eta_p^2 = .45$, with significantly greater dependency for closed- than open-loops. No other significant main effect or interaction was seen, Fs < 1.71, ps > .20. A 2x2 (Tested x Loop) ANOVA on dependency revealed a main effect of Loop, $F(1,25) = 14.21$, $p < .001$, $\eta_p^2 = .36$. No other significant effect or interaction was observed, Fs < 2.83, ps > .96.

2.6.3. Discussion

Experiment 4 replicated Experiments 1-3, showing consistent dependency for closed-loops despite high levels of forgetting. Experiment 4 showed higher retrieval accuracy for closed- than open-loops at T2, consistent with Experiment 3. This accuracy difference was present at T1 (consistent with results reported in Horner et al., 2015), however the effect was significantly greater after a week. This suggests that the structure of the overlapping associations can affect long-term retention, but seemingly only when the structures are manipulated in a within-subject design (as in Experiments 3-4, relative to Experiment 2).

2.7. A model of independent forgetting

Across four experiments I provide evidence for consistent levels of retrieval dependency, despite varying levels of forgetting. Evidence for dependency is consistently seen for closed-loops at both T1 and T2 and, importantly, dependency is retained despite variable levels of forgetting. If event-based representations, or closed-loops more specifically, fragment as a function of forgetting, such that some aspects of the memory trace are forgotten more quickly than others (e.g., *Kitchen* is forgotten, but not *hammer* or *Barack Obama*), then we would expect to see a decrease in dependency over time.

To ensure that the levels of dependency seen at T2 across Experiments 1-4 were greater than expected if forgetting was independent, I used a model of independent forgetting. The independent model used in the main analyses predicts the level of dependency if the retrieval of associations from an event are independent. It takes into account each participant's retrieval accuracy separately at each time point but does not take into account the amount of forgetting for each participant. Here, I used a model that predicted the level of dependency at T2, given a participant's retrieval accuracy and rate of forgetting between T1 and T2. The model can therefore be thought of as a model of independent forgetting, as opposed to a model of independent retrieval. The model predicts the level of dependency expected at T2 if closed-loops are forgotten in an independent manner.

Participants' retrieval data across all closed-loops and pairwise associations at T2 was simulated based on their performance at T1 and their overall level of forgetting (i.e., the difference in retrieval accuracy between T1 and T2). Note that the model includes a single forgetting parameter, such that the mean rate of forgetting is (on average) uniform across all closed-loops and element-types. The simulated data was based on the assumption of independent forgetting, such that the forgetting of one pairwise association for a closed-loop was not predictive of forgetting for any other pairwise association for that closed-loop. Specifically, I took performance for each T1 retrieval trial, across all cue-test pairs, which resulted in a 6x*N* matrix (with 6 cue-test pairs, and *N* events) where each trial was either correct or incorrect. I then simulated performance at T2 by converting correct trials to incorrect trials randomly until mean performance for the 6x*N* matrix was equated to the observed performance at T2. Importantly, each correct trial at T1 had a probability of being simulated as incorrect at T2 based on the mean level of forgetting for that participant. I then calculated the level of retrieval dependency for this simulated data set. This gives an estimate of the level of dependency for an individual participant at T2 under the assumption of independent forgetting. If the dependency seen in the observed data is greater than the simulated data, then this provides positive evidence that the forgetting of pairwise associations for closed-loops does not occur in an independent manner. For each participant, I simulated 100 data sets, and present the mean dependency across these simulations. Because Experiment 3 did not include an immediate test, I could only assess independent forgetting in Experiments 1, 2 and 4.

2.7.1. Results

Mean proportion of joint retrieval (and standard deviations) for the data and independent model at T2 for the simulated and observed data are presented in Table 7.

Table 7 Mean proportion of joint retrieval (and standard deviations) for the data and independent model at T2 for the simulated and observed data in Experiments 1, 2, and 4. For experiment 1, the proportion of joint retrieval is collapsed across the sleep and awake condition. Only trials where participants retrieved cue-target associations not previously tested at T1 are included.

	Simulated		**Observed**	
	Data	**Independent**	**Data**	**Independent**
Experiment 1	.68 (.16)	.65 (.17)	.69 (.15)	.66 (.17)
Experiment 2	.61 (.08)	.60 (.09)	.62 (.08)	.58 (.09)
Experiment 4	.58 (.08)	.57 (.07)	.60 (.09)	.57 (.07)

In order to assess whether the model predicts a decrease in dependency between T1 and T2 if forgetting is independent, I first assessed whether dependency for closed-loops in the simulated data at T2 showed a decrease relative to the observed dependency at T1. A 2x3 (Session x Experiment) mixed ANOVA (where the within-subject factor Session refers to the observed dependency at T1 vs the simulated dependency at T2 and the between-subject factor

Experiment refers to Experiment 1, 2 and 4) revealed a significant main effect of Session, $F(1,101) = 61.91$, $p < .001$, $\eta_p^2 = .38$, with dependency decreasing between T1 and T2. An interaction between Session and Experiment was also seen, $F(2,101) = 11.77$, $p < .001$, $\eta_p^2 = .19$, with a significantly greater decrease in dependency between T1 and T2 in Experiment 2, relative to Experiment 1, $t(25) = 3.52$, $p < .01$, $d = 1.09$, and Experiment 3, $t(25) = 2.18$, $p = .04$, $d = .71$ (note this effect did not survive correction for multiple comparisons). Further, a Bayesian pairwise *t*-test comparing the observed dependency at T1 with the simulated dependency at T2 across all participants from Experiments 1, 2 and 4 ($N = 104$) revealed a Bayes factor > 1000 in favour of the hypothesis that dependency should decrease as a function of independent forgetting (exceedance probability > .99; prior Cauchy distribution r = .707, centred at 0; H_0 = no decrease in dependency between T1 and T2), providing strong evidence that dependency should decrease if closed-loops fragment as a function of forgetting.

I next assessed whether dependency was greater for closed-loops in the observed data at T2 relative to the simulated data at T2. A 2x3 (Model x Experiment) ANOVA (with the within-subject factor Model referring to the observed dependency at T2 vs the simulated dependency at T2) revealed a significant main effect of Model, $F(1,101) = 10.43$, $p < .01$, $\eta_p^2 = .09$, with greater dependency in the observed relative to the simulated data. Consistent with this, a Bayesian pairwise *t*-test comparing the observed dependency at T1 with the observed dependency at T2 across all participants from Experiments 1, 2 and 4 ($N = 104$) revealed a Bayes factor = 5.41 in favour of the null hypothesis that dependency does not decrease as a function of forgetting (exceedance probability = .84; prior Cauchy distribution r = .707, centred at 0; H_0 = no decrease in dependency between T1 and T2), providing positive evidence that dependency does not decrease over time.

2.7.2. Discussion

The independent forgetting model estimates dependency at T2 under the assumption that cue-target associations are independently forgotten. Compared to the observed dependency, the simulated data shows that independent forgetting of individual event elements predicts a significant decrease in dependency between T1 and T2. Critically, I observed significantly greater dependency in the observed data at T2 relative to the simulated data. A Bayesian analysis also provided positive evidence that the observed dependency does not decrease as a function of forgetting for closed-loops.

2.8. General Discussion

Across four experiments I provided consistent evidence for retrieval dependency for closed-loops after a delay, despite variable levels of forgetting. Critically, comparisons with a model of independent forgetting provided evidence for levels of dependency for closed-loops that are greater than that predicted under an assumption of independent forgetting. This is consistent with the hypothesis that closed-loops tend to be forgotten in an all-or-none manner, being more likely to either be retained or forgotten in their entirety.

Experiment 1 showed that dependency for closed-loops is retained across a 12-hour interval, irrespective of whether encoding and immediate retrieval is followed by sleep or wakefulness. Consistent with previous findings (e.g., Barrett & Ekstrand, 1972; Gais, Lucas, & Born, 2006; Jenkins & Dallenbach, 1924; Lahl, Wispel, Willigens, & Pietrowsky, 2008; Plihal & Born, 1997; Tucker et al., 2006), I observed that sleep reduced forgetting, but did not change the form that forgetting takes. Experiment 1 also provided an opportunity to assess whether sleep (relative to wakefulness) supports the integration of overlapping information (Lau et al., 2010). Here I found no evidence to suggest that sleep promoted integration of open-loops, but it did appear to improve participants' ability to make inferences between related information (i.e., inferring that B-C elements are related after directly encoding A-B and A-C pairs).

In Experiment 2, the interval between encoding and retrieval was extended to one week. Despite numerical increases in forgetting in Experiment 2 relative to Experiment 1, I saw no evidence for changes in dependency for closed- or open-loops. Experiments 3 and 4 replicated Experiment 2, providing further evidence for dependency for closed-loops when the interval between encoding and retrieval was one week. Experiments 3 and 4 also showed that the associative structure of the underlying representation can support both immediate and long-term retention. Retrieval accuracy was higher for closed- than open-loops; a difference that increased significantly over the course of a week. Interestingly, this effect was only seen in a within-subject design where each participant learnt both closed- *and* open-loops (in Experiments 3 and 4), but not in a between-subject design where each participant either learnt closed- *or* open-loops (in Experiment 2).

Despite a long-standing interest in forgetting, little research has focussed on how memory representations change over a period of forgetting. Here I used the presence of dependency to infer the coherence of an underlying memory representation and asked how dependency changes as a function of forgetting. A similar approach was used by Brady et al. (2013) who assessed dependency for specific properties of an object (i.e., 'exemplar' and 'state'). Brady et al. (2013) found that dependency decreased over time, such that the exemplar (e.g., shape of a glass) and state (e.g., contents of the glass) of an object were forgotten independently. Thus, object- or item-based representations appear to fragment over time. This result appears inconsistent with the current results where no evidence for a decrease in dependency was seen for closed-loops over time. However, the experiments presented here were specifically interested in forgetting of coherent event-based representations (i.e., closed-loops) that require the binding of three distinct elements or items (i.e., a person, location, and object). Thus, whereas Brady et al. (2013) focussed on item-based representations; specifically multiple aspects of an individual item, the current experiments focussed on event-based representations; specifically associations between multiple distinct items/elements.

Neuropsychological and neuroimaging studies have demonstrated a functional dissociation between regions of the medial temporal lobe, with the perirhinal cortex supporting item-based representations and the hippocampus supporting event-based representations (Barense et al., 2005; Davachi et al., 2003; Diana, Yonelinas, & Ranganath, 2010; Lee et al., 2005; Ranganath et al., 2004). Although the perirhinal cortex has been implicated in certain associative processes (Mayes et al., 2007; Mayes et al., 2004) and conjunctive representations (Barense et al., 2007; Bussey & Saksida, 2007; Watson et al., 2012) of items/objects, such regions are not thought to support associations between multi-modal representations (Diana et al., 2007; Eichenbaum et al., 2007; Mayes et al., 2007). Instead, the hippocampus has been implicated in the multimodal binding that is required to form more complex event representations (Cohen et al., 1999; Damasio, 1989; Davachi, 2006; Eichenbaum et al., 2007; Horner et al., 2012). The imagery task and memoranda used here were designed to require cross-modal binding (Horner et al., 2015; Horner & Burgess, 2013, 2014) given that the hippocampus has been shown to act as a convergence zone (Backus et al., 2016); binding multi-modal information into coherent event representations (Damasio, 1989; Marr, 1971; Teyler & DiScenna, 1986). Thus, it is possible that the differences in forgetting seen between the present experiments and Brady et al. (2013) relate to a dissociation between item-based perirhinal representations and event-based hippocampal representations.

This dissociation is also apparent in the psychological literature in relation to retrieval dependency. Although Brady et al. (2013) saw dependency between the retrieval successes of exemplar and state information immediately after encoding, Brady et al. (2013) provided evidence that such dependency might be primarily driven by encoding-related factors. Horner & Burgess (2014) provided evidence against an encoding-based explanation of dependency for closed-loops, by separating out the encoding of pairwise associations for three element events, presenting each pairwise association in separate encoding trials (as in the present experiments). No difference in dependency was seen between this ‘separated’ encoding condition relative to when all three elements were encoded on a single trial (Horner & Burgess, 2014). This suggests that, even when tested immediately after encoding, the dependency seen for item-based and event-based representations might be driven by different factors. In the case of item-based representations; encoding-related factors such as attention, and in the case of event-based representations; perhaps a retrieval-related process that allows for holistic retrieval.

Evidence from the source-memory literature also suggests that there is a degree of asymmetry in relation to how source details are bound to items, with information about colour and location being directly, but independently bound to item information, but not each other (Starns & Hicks, 2005, 2008; see Hicks & Starns, 2015 for review). This lack of coherence and symmetry for item-based representations might underlie the decreases in dependency seen over time in Brady et al. (2013). As distinct aspects of an item are stored in a relatively independent manner, they are likely to also be forgotten in a similarly independent manner. In contrast, event-based representations are more coherent and symmetrical in nature (Horner et al., 2015; Horner & Burgess, 2013) and as such the forgetting of elements from the same event is more likely to be related. Interestingly, Sekeres et al. (2016) recently showed that ‘peripheral’ details from event-based memories are forgotten more rapidly than ‘central’ details. It is possible that this difference in forgetting rates for peripheral and central details might relate to the differences in the pattern of forgetting for event- and item-based representations observed here and in Brady et al (2013); with central details of an event (e.g., *Barack Obama* in the *kitchen* with a *hammer*) being more likely to be forgotten in an all-or-none manner, and peripheral/item-based details (e.g., the colour vs. shape of *Barack Obama’s* tie) being forgotten in an independent way.

It has recently been proposed that representations that rely on the perirhinal cortex are more likely to be forgotten as a result of interference (as opposed to decay) (Sadeh et al., 2014), as neural representations for items in the perirhinal cortex are likely distributed and overlapping (Norman & O'Reilly, 2003). Thus, it is possible that encoding similar objects results in interference due to their representational overlap. However, event-based representations are thought to be encoded in the hippocampus, where a pattern separation process supported by the dentate gyrus (Bakker et al., 2008; Berron et al., 2016; Leutgeb, Leutgeb, Moser, & Moser, 2007; Neunuebel & Knierim, 2014) and more sparse representations (Barnes, McNaughton, Mizumori, Leonard, & Lin, 1990; Viskontas, Knowlton, Steinmetz, & Fried, 2006) are likely to reduce representational overlap between similar events. This decrease in representational overlap decreases the likelihood of interference (McClelland et al., 1995). Instead, forgetting for hippocampal representations is thought to be a result of decay (Hardt et al., 2013). Sadeh et al. (2016) provided behavioural evidence for this proposed dissociation, showing that whereas recollection (a process thought to be supported by hippocampal representations) decreased as a function of time between encoding and retrieval (consistent with forgetting via decay), familiarity (a process thought to be supported by perirhinal representations) decreased as a function of experimentally induced interference. Recent evidence using 'precision' measures for assessing object-colour memory has shown that encoding similarly coloured objects causes interference, leading to a decrease in precision (i.e., participants can still remember the colour, but with less specificity than previously) (Sun et al., 2017), supporting the notion that forgetting for item-based representations is more likely driven by interference than decay.

The lack of a decrease in dependency can be taken as support for the hypothesis that coherent (closed-loop) events tend to be forgotten in an all-or-none manner. If hippocampal representations were forgotten by a process of decay, as proposed by Sadeh et al. (2014), this would suggest that decay is relatively uniform for the separate elements of an event. Although there is variation in the amount of decay *across* events, there would be less variation *within* an event. An alternative decay account would predict that variation is present in the rate of decay within an event, however the process of retrieval compensates for this variation. It has previously been shown that closed-loops are supported by the hippocampus (Horner et al., 2015), consistent with the proposal that event-based representations are retrieved by a process of pattern completion (see Horner & Doeller, 2017 for review). Pattern completion allows for the retrieval of a complete memory trace (i.e., pattern) in the presence of a partial or ambiguous input (Gardner-Medwin, 1976; Hopfield, 1982; Marr, 1971; McClelland et al., 1995; Treves & Rolls, 1992, 1994). Here, activation of a single event element (e.g., a location) triggers the reactivation of all other elements for that event (e.g., a person and an object). Under such an account, decay could be non-uniform within an event, but pattern completion in hippocampal subfield CA3 (Hopfield, 1982; Treves & Rolls, 1992), or more widespread recurrency within the hippocampal complex (Kumaran & McClelland, 2012), would lead to the presence of dependency as long as the associations between some elements are sufficiently strong. However, when decay is sufficient the remaining associations may no longer be able to support retrieval, entailing that the entire representation cannot be retrieved. In other words, although decay might be non-uniform within an event, pattern completion produces the appearance of uniformity at retrieval, leading to behavioural dependency. This same mechanism has been used to account for retrieval generalisation on paired-associate tasks (e.g., where participants make inference judgments about two overlapping associations; Banino et al., 2016).

While the precise mechanism that underlies the pattern of forgetting seen in the present experiments is unclear, the results (taken alongside those of Brady et al. (2013)) suggest that the forgetting of event-based representations is driven by a different mechanism than for item-based representations. It is possible that this is hippocampal-based and is supported by the known recurrent circuitry in this region that support the computational process of pattern completion. Further, this may be a result of decay, perhaps driven by neurogenesis of hippocampal granule cells (Frankland et al., 2013), or more active regulatory changes (Hardt et al., 2013). However, the findings presented here do not necessitate that forgetting occurs due to the erasure of the memory. It is possible that aspects of some events, or entire events, can spontaneously recover and be brought to mind at some later point (Tulving & Pearlstone, 1966), or could even be recovered via optogenetic induction (Roy et al., 2016; Ryan, Roy, Pignatelli, Arons, & Tonegawa, 2015).

Sleep is thought to not only play an active role in the strengthening of memory representations, but also in the integration of overlapping information (Lewis & Durrant, 2011; Stickgold & Walker, 2013). For example, Lau et al. (2010) presented evidence for an increase in participants' ability to infer the relationship between B-C pairs following a nap after directly encoding A-B and A-C pairs. This increase in inference ability following sleep was taken as evidence for sleep playing an active role in the integration of A-B and A-C pairs; similar to evidence for integration seen during repeated presentations of such pairs (Zeithamova et al., 2012). However, it has been suggested that B-C inference can readily be supported by retrieval-related processes, supported by the recurrent connections in the hippocampus (Kumaran & McClelland, 2012) where the relationship between B and C is generated 'on-the-fly' via the retrieval of A-B and then A-C. Under this account, the probability of successful inference increases via increases in the associative strength of the directly encoded A-B and A-C pairs. Here I used retrieval dependency to distinguish between these two accounts. In Experiment 1, participants learnt open-loops (A-B, A-C associative structures) and were, after sleep or wakefulness, tested on the directly encoded pairs and the B-C non-encoded pairs. Consistent with Lau et al. (2010), I found evidence that sleep increased performance on an B-C inference task (though the effect was relatively small). However, there was no evidence for increases in dependency as a function of sleep for open-loops. Similarly, no evidence for increases in dependency was seen for open-loops after a one-week delay (Experiments 2-4), that included multiple sleep-wake cycles. If sleep does play a role in mnemonic inference, this effect is likely to be primarily driven by increases in the associative strength for directly encoded pairs that allows for inference at the point of retrieval, rather than a more active sleep-related integration process.

For both open- and closed-loops, sleep decreased forgetting relative to wakefulness across a 12-hour delay (Experiment 1). Sleep appears to decrease forgetting but does not change the form that forgetting takes. This is consistent with existing models of systems consolidation (Frankland & Bontempi, 2005; McClelland et al., 1995; Squire & Alvarez, 1995) where sleep is thought to reduce forgetting by stabilising existing connections in the neocortex, perhaps counteracting memory decay within the hippocampus (Frankland et al., 2013). However, this process appears to occur without altering the form that forgetting takes. Interestingly, I also observed that prior testing diminished the effect of sleep on memory, in line with previous findings showing that retrieval practice may mitigate the effect of sleep on retention (Bäuml, Holterman, & Abel, 2014).

In Experiments 3 and 4 greater retrieval accuracy was seen for closed- than open-loops, with the difference in retrieval accuracy increasing over time. Thus, the associative structure formed at encoding appears to significantly modulate the amount of forgetting over the course of a week. No difference in retrieval accuracy between closed- and open-loops was seen after a week in Experiment 2. The key difference between these experiments is the between- vs within-subject experimental design. Whereas in Experiment 2 participants either learnt closed- *or* open-loops, in Experiments 3 and 4 participants learnt both closed- *and* open-loops. Importantly, overall retrieval accuracy after a week (averaged across closed- and open-loops) in Experiments 3 (.43) and 4 (.44) was similar to Experiment 2 (.43). Thus, overall forgetting rates were comparable, but learning both closed- *and* open-loops appeared to bias forgetting, such that open-loops were more likely to be forgotten than closed-loops.

One possible explanation for this effect is a competitive model of forgetting, where multiple memory representations compete for ‘survival’. Closed-loops are already associated with significantly higher retrieval accuracy at immediate test. If this higher retrieval accuracy relates to greater mnemonic strength, then perhaps this allows closed-loops to ‘out-compete’ open-loops to survive. Note though that Poulton (1982) argued that within-subject designs, in contrast to between-subject designs, can bias performance (in unknown ways) due to asymmetries in encoding and/or retrieval strategies when conditions are interleaved randomly. Accordingly, the difference in retrieval accuracy between closed- and open-loops observed here could be due to the transfer of a particular strategy that is appropriate in one condition (i.e., closed-loops), but not in the other (i.e., open-loops). Further work is needed to determine the precise experimental conditions under which closed-loops are more likely to be retained.

2.9. Conclusion

Across four experiments I provided consistent evidence that retrieval dependency does not change over time, despite variation in the overall amount of forgetting. This suggests that the associative structure formed at encoding has a consistent and lasting impact on the coherence of retrieval. In relation to closed-loops, the results support the hypothesis that coherent event representations tend to be forgotten in an all-or-none manner, such that events are more likely to be either forgotten or retained in their entirety. I also saw evidence that the associative structure formed at encoding can, under specific conditions, modulate the overall amount of forgetting. When participants learnt closed- *and* open-loops, forgetting rates were significantly lower for closed-loops (this was not the case when participants learnt either closed- *or* open-loops), suggesting that both retrieval accuracy and dependency can be modulated by the associative structure formed at encoding.

3. TMR-induced forgetting of open-loops during sleep

3.1. Abstract

Research suggests that memory reactivation during sleep may strengthen the association between overlapping memories and promote mnemonic integration. However, memory reactivation during sleep has also been associated with forgetting, depending on the interval between the encoding of two overlapping associations. Here, I assessed whether retrieval accuracy and dependency for open-loops - consisting of two separately encoded but overlapping pairwise associations - changed as a function of targeted memory reactivation (TMR) during sleep. No evidence for dependency was seen for open-loops, neither immediately after encoding nor following a period of sleep that included TMR. Instead, I observed evidence that TMR decreased retention for the second encoded association of an open-loop. This effect was only seen for open-loops not tested prior to sleep. These findings suggest that memory reactivation during sleep can lead to forgetting; possibly as an effect of proactive interference at encoding, but this is dependent on whether a memory is tested prior to sleep or not.

3.2. Introduction

Memory consolidation is the process where an initially labile memory is transformed into a relatively stable, long-term form (McGaugh, 1966). Consolidation is typically separated into synaptic and systems consolidation. Synaptic consolidation refers to post-encoding changes on a synaptic and cellular level that draw to a close within hours after memory acquisition. Systems consolidation, on the other hand, refers to a more protracted reorganisation of memory representations in the neocortex. This process may occur over days, weeks, months, or even years (see Dudai et al., 2015 for review). Theories of systems consolidation (e.g., Marr, 1971; McClelland et al., 1995; Squire & Alvarez, 1995) suggest that the repeated neural reinstatement of recently formed memory representations during sleep constitutes one core mechanism of systems consolidation. This proposal is supported by evidence from rodent studies showing that patterns of neural activity associated with spatial information acquired during learning are spontaneously reactivated during sleep (Wilson & McNaughton, 1994), with the amount of reactivation predicting later memory performance (Dupret et al., 2010; van de Ven, Trouche, McNamara, Allen, & Dupret, 2016). Similarly, neuroimaging findings in humans have shown that memory reactivation of novel information during sleep predicts subsequent memory performance (Deuker et al., 2013; Peigneux et al., 2004; see also Schapiro, McDevitt, Rogers, Mednick, & Norman, 2018; Staresina, Alink, Kriegeskorte, & Henson, 2013 for evidence for reinstatement during wakeful rest).

It has been proposed that systems consolidation during sleep strengthens the association between overlapping memory representations (Diekelmann & Born, 2010; Lewis & Durrant, 2011; Stickgold & Walker, 2013). This proposal is supported by evidence showing that participants are able to gain insight into logical problems (Wagner et al., 2004), and to infer the relationships between objects embedded in a hierarchical structure following sleep (Ellenbogen et al., 2007). This has been taken as evidence that sleep may support the formation of novel associations between overlapping, but otherwise orthogonal memory representations (Diekelmann & Born, 2010). Consistent with this, Lau et al. (2010) found that after learning overlapping A-B and A-C pairs, the ability to infer the relationship between B and C was increased for participants that slept, relative to those who remained awake between encoding and retrieval. This increase in B-C inference following sleep was interpreted as reflecting a process where A-B and A-C associations are integrated into a generalised representation during sleep (Lau et al., 2010), potentially forming a direct association between B and C (Shohamy & Wagner, 2008). However, inferring the relationship between B and C can also occur 'on-the-fly' via the retrieval of A-B and A-C, in the absence of a generalised representation (Banino et al., 2016; Kumaran & McClelland, 2012).

In Chapter 2, I attempted to distinguish between these two possibilities by assessing retrieval dependency for pairwise associations encoded as open-loops. Retrieval dependency has previously been used to show that the separate encoding of three overlapping associations between, for instance, *kitchen-hammer*, *kitchen-Barack Obama*, and *hammer-Barack Obama* in a closed-loop leads to the formation of a coherent memory representation (Horner et al., 2015; Horner & Burgess, 2014), similar to those formed during a single encoding trial (Horner & Burgess, 2014). However, no dependency is seen when only two out of the three possible pairwise associations are encoded (e.g., *kitchen-hammer*, *kitchen-Barack Obama*, but not *hammer-Barack Obama*) in an open-loop (Horner & Burgess, 2014) (analogous to the A-B, A-C associative structure used in Lau et al. (2010)). In Chapter 2, I assessed retrieval accuracy and dependency for open-loops immediately following encoding and following a 12-hour delay (that included either a period of sleep or wakefulness), under the assumption that if A-B and A-C associations are integrated during sleep, then dependency should be seen for open-loops following sleep. Thus, if increases in B-C inference were seen following sleep (as in Lau et al., 2010), with an associated increase in dependency, then this would provide evidence for the proposal that sleep supports inference via a generalised representation (Lau et al., 2010).

However, if increases in B-C inference were seen without increases in dependency, then this would support the proposal that the relationship between B-C can be inferred ‘on-the-fly’ at the point of retrieval (Kumaran & McClelland, 2012). I showed that sleep, relative to wakefulness, improved participants’ ability to infer the relationship between B-C pairs (consistent with Lau et al., 2010). However, no evidence for dependency was observed following sleep, suggesting that the role of sleep in increasing inference is unlikely to be driven by the formation of a generalised representation.

However, a distinction can be made between a passive role of sleep in protecting information from interference, and active systems consolidation processes associated with the neural reinstatement of mnemonic representations during sleep (see Ellenbogen, Payne, & Stickgold, 2006 for review). It is possible that memory reactivation during sleep is a necessary condition for the emergence of a generalised representation consistent with that proposed by Lau et al. (2010). This would be consistent with evidence for integration following repeated presentations of A-B and A-C pairs during wakefulness (Zeithamova et al., 2012). Further, recent work has shown that memory reactivation during both wakeful rest and sleep prioritises weakly encoded information, but with neural reactivations of weakly encoded information predicting memory improvements in participants who slept and not those who remained awake during a rest interval (Schapiro et al., 2018). Given these findings it is possible that memory reactivation during sleep may provide the optimal state for strengthening the association between the third, and final non-encoded pair for open-loops; resulting in retrieval dependency. This would be consistent with the proposal that systems consolidation during sleep supports a process of mnemonic integration for separately encoded A-B and A-C associations (Lau et al., 2010), potentially forming a direct association for the non-encoded B-C pair (Shohamy & Wagner, 2008).

Recently, studies have begun to use a technique called targeted memory reactivation (TMR) to assess the role of memory reactivation during sleep (see Oudiette & Paller, 2013 for review).

An early demonstration of this technique was reported by Rasch, Büchel, Gais, and Born (2007) who showed that participants' retention for multiple object-location pairs, associated with an odour, was increased when participants were re-exposed to the odour, relative to participants who were exposed to an odourless stimulus, during sleep. This initial demonstration provided evidence that the reapplication of an odour during sleep could act as a 'context' cue for the reactivation of memories associated with the odour (Rasch et al., 2007). More recently, studies have begun to leverage the use of auditory stimuli to target specific memory representations during sleep. For instance, Rudoy et al., (2009) had participants learn a series of object-location pairs with each object being presented alongside a characteristic sound (e.g., object: *dog* - sound: *'bark'*). Rudoy et al., (2009) showed that retention for object-location pairs associated with sounds presented during sleep was improved, relative to pairs associated with sounds not presented during sleep. Following these initial findings, evidence for a positive effect of TMR during sleep on retention has accumulated (see Cellini & Capuozzo, 2018 for review).

However, previous work has tended to focus on the effect of TMR during sleep on the retention of single associative memories (but see Antony, Cheng, Brooks, Paller, & Norman, 2018; Oyarzún, Morís, Luque, de Diego-Balaguer, & Fuentemilla, 2017; Payne, Stickgold, Swanberg, & Kensinger, 2008). As systems consolidation during sleep is thought to be selective (Oudiette , Antony, Creery, & Paller, 2013; Rauchs et al., 2011; Saletin, Goldstein, & Walker, 2011), prioritising the consolidation of some memories over others (Cairney, Lindsay, Sobczak, Paller, & Gaskell, 2016; Drosopoulos, Schulze, Fischer, & Born, 2007; Fischer & Born, 2009; Schapiro et al., 2018; Schmidt et al., 2006; Wilhelm et al., 2011), it is less clear how TMR during sleep may affect memory for two or more overlapping associations.

Oyarzún et al. (2017) recently demonstrated that TMR of overlapping memories during sleep can lead to forgetting. In this study, the encoding of overlapping A-B (i.e., an object presented at a spatial location on a spatial grid) and A-C (i.e., the same object presented at a different location on a spatial grid) pairs was manipulated such that the encoding of A-C followed either immediately after encoding the A-B pair, or after a 3-hour delay. Oyarzún et al. (2017) reasoned that if memory strength decays as a function of time, then increasing the interval between the encoding of two overlapping associations should increase the likelihood of retroactive interference. As such, C, relative to B, should be more strongly associated with A when the encoding of A-C followed a 3-hour delay. This was not assumed to be the case when A-B and A-C pairs were encoded immediately after each other (Oyarzún et al., 2017). Oyarzún et al. (2017) showed that re-exposing participants to an auditory stimulus (associated with the encoding of A-C) during sleep, decreased retention for the A-B pair when encoding was separated by a 3-hour delay. This suggests that memory reactivation during sleep can lead to forgetting. Interestingly, when the overlapping associations were encoded immediately after each other, TMR supported retention for the A-B pair. However, this study did not assess memory performance for the A-C association. This point is critical given evidence that repeated retrievals of a 'target' memory during wakefulness can improve retention for the 'target', but also induce long-lasting decreases in performance for a 'competing' memory representation (Anderson et al., 1994); a phenomenon referred to as retrieval-induced forgetting (see Anderson, 2003 for review).

Anderson et al. (1994) showed that if participants learnt six category exemplars for fruits, repeated retrievals of three of the six exemplars in response to an exemplar stem (e.g., *fruit: or_* for the exemplar *orange*) decreased memory performance for exemplars (e.g., *banana*) not retrieved prior to a critical memory test. More recently, Wimber et al. (2015) used fMRI to demonstrate that the repeated retrieval of an A-B pair led to the 'suppression' of cortical representations associated with a competing A-C pair. Wimber et al. (2015) also showed that the level of neural suppression induced by repeated retrievals predicted forgetting for A-C (but see Potter et al., 2018). If repeated retrievals during wakefulness can lead to forgetting for ‘competing’ memory representations (Anderson et al., 1994; Wimber et al., 2015), then it is possible that benefits in retention for an A-B association, as reported by Oyarzún et al. (2017), may also be associated with decreases in retention for the overlapping A-C association. This would be consistent with recent work showing that TMR can lead to increases or decreases in retention for two overlapping associations depending on which association was paired with high or low future reward (Antony et al., 2018).

Here, I used a similar design to that developed in Horner and Burgess (2014) (see Figure 5). At encoding, participants learned a series of triplets, each consisting of three elements (famous people, locations, and common objects). Triplets were ‘built up’ over two separate encoding trials with each trial consisting of one of the three possible pairwise associations for each triplet. In contrast to Chapter 2, here encoding only allowed for forming open-loops*;* where only two out of the three possible pairwise associations are encoded. As the encoding of overlapping material is susceptible to proactive interference (see Postman, 1971 for review), it is possible that the first and second encoded associations for each open-loop differ in their mnemonic strength (similar to that seen in Oyarzún et al., 2017). During encoding, each of the pairwise associations belonging to the same open-loop were associated with a characteristic auditory stimulus that was presented alongside the word-pairs. For instance, if the participant was presented with the associations *Barack Obama-hammer* and *kitchen-hammer*, the object word (i.e., ‘*hammer*’) would be presented acoustically alongside both word-pair presentations. At both immediate and delayed retrieval, the learnt associations were tested in both directions (e.g., cue: *hammer*, target: *Barack Obama*; cue: *Barack Obama*, target: *hammer*) using six-alternative forced choice among elements of the same category from different open-loops. At immediate retrieval, pairwise associations belonging to 30 out of the 60 open-loops were tested in order to control for possible effects of prior testing. Associations from all 60 open-loops were tested at delayed retrieval.

During sleep, between immediate and delayed retrieval, I presented auditory stimuli associated with specific open-loops to participants in order to induce memory reactivation of the learnt associations. Following the findings presented in Chapter 2, I first assessed whether TMR during sleep (rather than sleep more generally as in Chapter 2) increased retention for open-loops. If memory reactivation during sleep plays an active role in strengthening the association between overlapping mnemonic representations (Diekelmann & Born, 2010; Lewis & Durrant, 2011), TMR during sleep should not only improve retention, but also lead to increases in dependency for open-loops. This would be consistent with the proposal that systems consolidation during sleep supports the formation of a generalised representation between two overlapping, but separately encoded associations (Lau et al., 2010). However, in light of evidence that TMR during sleep can in some instances lead to forgetting (Antony et al., 2018; Oyarzún et al., 2017; Simon, Gómez, & Nadel, 2018), I also assessed whether TMR during sleep may affect retention for open-loop associations differently; possibly as an effect of proactive interference at encoding, such that TMR supports memory for the first encoded association but decreases retention for the second encoded association of an open-loop.

3.3. Methods

3.3.1. Participants

As previous work has primarily reported positive effects of TMR on retention, an effect size *d* = .40 for detecting an increase in retention following TMR was calculated from previous published work assessing memory for single associations immediately following sleep (Ashton, Cairney, & Gaskell, 2018; Cairney, Guttesen, El Marj, & Staresina, 2018; Cairney et al., 2016; Creery, Oudiette, Antony, & Paller, 2015; Göldi, van Poppel, Rasch, & Schreiner, 2019; Oudiette et al., 2013; Rudoy et al., 2009; Schreiner & Rasch, 2015; Van Dongen et al., 2012), with *n* = 214. Cohen's *d* was calculated from reported *t*-statistics divided by the square root of the within-subject sample size (Lenth, 2006). Using G* Power (Faul et al., 2009), with $d = .40$, I estimated that a sample size of $n = 37$ was required to detect a significant positive effect of TMR, if one is present, at a power of .65 and $\alpha = .05$.

A total of thirty-nine participants took part in the experiment. Participants were recruited from the University of York student population and took part in exchange for course credit or monetary compensation. Written informed consent was obtained from all participants. Participants were informed that they would be taking part in a memory study but were not made aware that TMR would be taking place during the sleep period. After completing the experiment, participants were asked if they had been aware of the auditory stimulation during the sleep period. Ten participants were excluded for the following reasons: awareness of TMR during sleep ($n = 2$), less than two full rounds of TMR administered during sleep ($n = 7$), and performance not exceeding chance (< ~16.7%) during the first test session ($n = 1$). Therefore, the analyses reported here were carried out on 29 participants (17 male, $M \pm SD$ age = 20.97 ± 4.02).

The reduction in the sample size included in the main analyses led to a decrease in power from .65 to .55 with which a positive effect of TMR, if one is present, can be detected with $\alpha = .05$ and $d = .40$. Note though that 8 out of the 10 studies included in the estimation of the effect size have assessed the effect of TMR across a nap session lasting ~90 mins, rather than overnight sleep (as in the current experiment). This will impact the total duration of sleep and full rounds of TMR that can be administered. This may lead to an underestimation of the power with which a positive effect of TMR can be detected in the current experiment. Note also that the average sample size of the studies used for estimating the effect size was 21 (range = 12-30) as compared to 29 in the current experiment.

The experiment was approved by the Research Ethics Committee of the Department of Psychology, University of York.

3.3.2. Materials

The stimuli consisted of 60 locations (e.g., *kitchen*), 60 famous people (e.g., *Barack Obama*) and 60 common objects (e.g., *hammer*). From these, 60 randomised location-person-object open-loops were generated for each participant. The auditory stimuli consisted of 60 spoken object words. The spoken object words corresponded to the 60 common objects used in the main stimuli. Spoken object words, rather than object sounds, corresponding to the main object stimuli were used to ensure congruency between the two stimulus sets. All object sounds (*M* ± *SD* duration = 656.17 ± 153.34-ms) were recorded in a female voice and adjusted for perceived loudness using Adobe Audition (Version 3.0).

3.3.3. Procedure

The experiment consisted of a single encoding session, a single sleep session, two retrieval sessions and a single inference session. Session 1 took place at ~9.30 pm (*M* = 9:36 pm, range = 8:13-10:18 pm). Session 1 consisted of a single encoding phase and an immediate retrieval phase (T1) (see details below). Session 2 took place at ~7.30 am (*M* = 7:27 am, range = 7:04-7:54 am) immediately following sleep. Session 2 consisted of one retrieval (T2) and one inference phase (see details below). The Pittsburgh Sleep Quality Index (Buysse, Reynolds, Monk, Berman, & Kupfer, 1989) was administered prior to the encoding phase. This was done in order to assess for any difficulties in sleep up to a month prior to the participant taking part in the experiment. Sleep quality ratings indicated that participants had no current history of sleep difficulties (*M* ± *SD* = 4.55 ± 1.80 *SD*). Self-report ratings of sleepiness were also collected prior to encoding (*M* ± *SD* = 3.07 ± 1.18) and second test phase (*M* ± *SD* = 2.50 ± 0.92) using the Stanford Sleepiness Scale (Hoddes et al., 1973). Ratings of sleepiness were not collected from one participant. There was no significant difference in self-reported ratings of sleepiness before (*M* ± *SD* = 3.07 ± 1.18) and after sleep (*M* ± *SD* = 2.50 ± 0.92), $t(27) = 1.95$, $p = .06$, $d = .37$ (but note the marginal *p*-value).

Encoding. During encoding, participants were presented with specific pairwise associations for each of the 60 open-loops. Participants learnt one pairwise association per trial. All pairwise associations were presented on a computer screen as words, with one item to the left and one to the right of fixation. The left/right assignment was randomly chosen on each trial. The words remained on screen for 6s. Participants were instructed to imagine, as vividly as possible, the items interacting in a meaningful way for the full 6s. Each word-pair presentation was preceded by a 500-ms fixation and followed by a 500-ms blank screen.

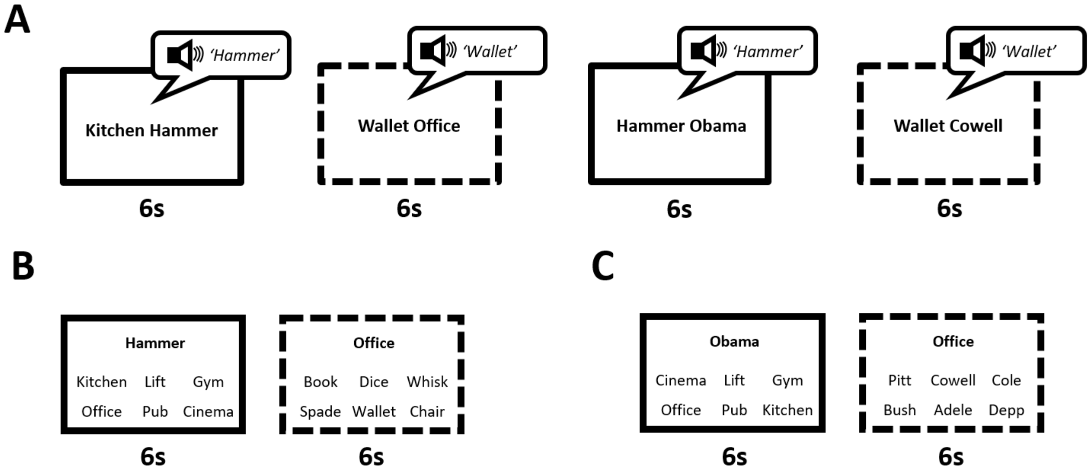

Figure 5 Experimental design. **(A)** Encoding. Participants saw multiple pairwise associations. They imagined each association 'interacting in a meaningful way as vividly as possible' for 6s. Each association was preceded by a 500-ms fixation and followed by a 500-ms blank screen. Solid and dotted lines were not presented but highlight pairwise associations belonging to separate open-loops. An event-specific auditory stimulus was presented alongside each pairwise association 1s after trial onset and 1s prior to the trial offset. **(B)** Retrieval. Participants were presented with a single cue and required to retrieve one of the other elements from the same open-loop from among five foils (elements of the same type from other open-loops) in 6s. Each cued-recognition trial was preceded by a 500-ms fixation and followed by a 500-ms blank screen. **(C)** Inference. Participants were presented with a single cue and required to infer which one of the elements was from the same open-loop from among five foils from different open-loops.

The encoding phase consisted of two blocks of 60 trials with one pairwise association from each open-loop presented during each block, making a total of 120 encoding trials. A break of 20s would follow every 30 encoding trials. Within each block, the order of presentation was randomised. Each open-loop consisted of a common object (e.g., participants learnt object-person, and then object-location). The encoding order for the open-loops across the two blocks was either (1) object-person, object-location or (2) object-location, object-person. The encoding order was pseudo-randomised for each participant, such that for each participant the encoding-order for 30 of the open-loops was object-person, object-location, and for the other 30 open-loops the encoding-order was object-location, object-person. On each encoding trial, the spoken object word was presented twice acoustically. For instance, if the participant was presented with the association *Barack Obama-hammer*, the spoken object word (i.e., ‘*hammer*’) would be presented alongside the word-pair presentation 1s after trial onset and 1s before the trial offset.

Immediate retrieval (T1). During retrieval, participants performed a six-alternative forced-choice cued-recognition task. On a given trial, the cue and six possible targets were presented simultaneously on screen. The cue was presented in the middle of the screen with the six possible targets; one target and five foils from the same category (e.g., if the target was *hammer*, the five foils would be other randomly selected objects from other open-loops), presented in two rows of three below the cue. The position of the target was randomly selected on each trial. Participants had 6s to respond with a key press. Each retrieval trial was preceded by a 500-ms fixation and followed by a 500-ms blank screen. For T1, 30 out of the 60 open-loops were tested. Each open-loop was tested with one of the cue-target associations (e.g., cue: object, target: location) in both directions (e.g., cue: object, target: location; cue: location, target: object). The cue-target associations were randomly presented across four blocks (with a single pairwise association from each open-loop tested in each block), making a total of 120 trials. The presentation order was randomised for each block. A 20s break would follow every 30 trials. Missing responses ($M \pm SD = 6.35 \pm 5.40\%$) were treated as incorrect trials

TMR set up. The 30 open-loops tested at T1 were rank-ordered by retrieval accuracy (i.e., mean performance per open-loop) in a descending order. Of the rank-ordered open-loops, 15 open-loops were assigned in an interleaved manner (from most to least accurate) to the TMR condition, and the remaining 15 were assigned to the non-TMR condition. Of the 30 open-loops not tested at T1, 15 open-loops were randomly allocated to the TMR condition, and the remaining 15 open-loops were allocated to the non-TMR condition. Accordingly, (1) 15 open-loops tested at T1 were assigned to the TMR condition, (2) 15 open-loops not tested at T1 were assigned to the TMR condition, (3) 15 open-loops tested at T1 were assigned to the non-TMR condition, and (4) 15 open-loops not tested at T1 were assigned to the non-TMR condition. The spoken object words belonging to open-loops assigned to the TMR condition were presented to participants during the TMR phase. An additional set of 6 control object words, that participants had not encountered at encoding, were intermixed with the main auditory stimuli.

Sleep period. The sleep period began at ~11 pm and lasted ~8 hours (*M* ± *SD* = 7 hrs, 55 mins ± 21 mins). Participants were left to sleep in a laboratory bedroom while their brain activity was monitored with polysomnography. During sleep, the auditory stimuli were presented via a speaker - mounted ~1.5 m above the bed - that was connected to an amplifier in a separate control room. Given the purported role of slow-wave sleep in systems consolidation (Born, 2010; Diekelmann & Born, 2010), TMR was initiated when participants were in non-rapid eye movement (NREM), sleep stage N3. The auditory cues were presented in a randomised order at a sound intensity of ~40dB (as measured with a sound-level meter placed at the same position where participants laid their heads on the pillow). In order to avoid any auditory feedback that may occur when sound cues are presented too rapidly in succession of each other (Schreiner, Lehmann, & Rasch, 2015), each sound presentation was separated by a 5s (with *M* ± *SD* = 0 ± 300-ms random jitter) interval following the sound offset. TMR continued for as long as participants were in sleep stage N3 between sleep onset and ~4 am (*M* ± *SD* number of full rounds of TMR = 7.34 ± 5.20). TMR was immediately paused if participants showed signs of arousal or awakening, or moved into sleep stages N1 or N2, or rapid eye movement (REM) sleep.

An Embla N7000 PSG system, with RemLogic (Version 3.4) software, was used to monitor sleep. EEG scalp electrodes were attached according to the international 10-20 system at 8 locations: frontal (F3, F4), central (C3, C4), parietal (P3, P4), and occipital (O1, O2), and each referenced to electrodes on the contralateral mastoid (A1 or A2). Left and right electrooculography electrodes were also attached, as were electromyography electrodes at the mentalis and submentalis bilaterally. A ground electrode was attached to the forehead. Impedance for all electrodes was set to < 5 k**Ω** prior to the sleep period. Online signals were unfiltered and digitally sampled at 200 Hz. All electrode sites were cleaned with NuPrep exfoliating agent before gold plated electrodes were attached using EC2 electrode cream. Offline sleep scoring was carried out in accordance with the criteria of the American Academy of Sleep Medicine (Iber & Iber, 2007) by two independent scorers. One participant's sleep data was not scored due to a recording error during the sleep period. Mean durations (and standard deviations) in minutes of sleep and percentage duration of sleep stages N1, N2, N3 and REM from both scorers are presented in Table 8.

Table 8 Mean duration (and standard deviations) of total sleep time in minutes, and N1, N2, N3 and REM in percentages as scored by two independent scorers.

	Total sleep (mins)	**REM (%)**	**N1 (%)**	**N2 (%)**	**N3 (%)**
Scorer 1	413.08 (39.36)	17.68 (4.26)	9.48 (5.18)	49.41 (6.99)	23.08 (7.02)
Scorer 2	414.50 (39.44)	23.35 (5.01)	8.05 (6.74)	43.04 (9.60)	25.55 (6.91)

Delayed retrieval (T2). During T2, participants performed the same cued-recognition task as during T1 but with all 60 open-loops now tested, making a total of 240 trials. The presentation order was randomised for each of the four retrieval blocks. A 20s break followed every 30 trials. Missing responses (*M* ± *SD* = 4.97 ± 4.93%) were treated as incorrect trials. Immediately following T2, participants performed an additional inference task. For example, if a participant had encoded the associations *Barack Obama-hammer* and *hammer-kitchen*, the non-encoded association between *Barack Obama* and *kitchen* would be tested during the inference task. For the inference task, the non-encoded associations for each open-loop were tested in each direction across two blocks, making a total of 120 trials. Participants had 6s to respond with a key press. Missing responses (*M* ± *SD* = 6.32 ± 7.29%) were treated as incorrect trials. Note that in contrast to Chapter 2, participants were explicitly instructed that they were being tested on the non-encoded associations.

3.3.4. Statistical analyses

For the analysis of retrieval accuracy, a 2x2 (Session x TMR) ANOVA is reported, with the within-subject factor Session referring to performance at T1 vs T2 and TMR referring to open-loops in the TMR vs non-TMR conditions. For T2, only open-loops not previously tested at T1 are included in the ANOVA. A further a 2x2 (Tested x TMR) ANOVA for retrieval accuracy at T2 is also reported, with the within-subject factor Tested referring to whether open-loops had previously been tested at T1 or not. For the dependency analysis (see Chapter 2 for a description of how retrieval dependency is modelled), a 2x2 (Session x TMR) ANOVA and a 2x2 (Tested x TMR) are reported. Here the dependent variable refers to the difference in the proportion of joint retrieval in the data and independent model.

For inference performance, I report a 2x2 (Tested x TMR) ANOVA with the factor Tested now referring to whether the non-encoded pairs were associated with open-loops previously tested or not tested at T1. Lastly, I report a 2x2x2 (Session x TMR x Encoding-order) ANOVA and 2x2x2 (Tested x TMR x Encoding-order) ANOVA to assess for differences in retrieval accuracy across open-loop associations as an effect of encoding-order (i.e., first vs second encoded associations).

Alpha was set to .05 (two-tailed) for all statistical tests. For each ANOVA, I report a partial eta-squared effect size (η_p^2). For *t*-tests, a Cohen's *d* is reported as the mean difference between conditions divided by the pooled standard deviation across conditions (Lakens, 2013). All statistical analyses were conducted using JASP (jasp-stats.org; JASP Team, 2018). For each *t*-test, a Bayes factor in favour of the alternative hypothesis is also reported. Bayes factors were computed with a Cauchy prior centred at 0 and a distribution parameter $r = .707$

Note that because open-loops tested at T1 were assigned to the TMR and non-TMR conditions by rank-ordering performance in a descending order, retrieval accuracy at T1 for open-loops in the TMR vs non-TMR conditions differed. In order to equate retrieval accuracy at T1 for open-loops in the TMR vs non-TMR conditions, (1) the lowest performing open-loop in the non-TMR condition and (2) the open-loop closest to mean performance at T1 in the TMR condition for each participant were removed from the analyses. Accordingly, for all analyses, the number of open-loops per condition was (1) 14 open-loops tested at T1 in the TMR condition, (2) 14 open-loops tested at T1 in the non-TMR condition, (3) 15 open-loops not tested at T1 in the TMR condition, and (4) 15 open-loops not tested at T1 in the non-TMR condition.

3.4. Results

3.4.1. Retrieval accuracy

Mean proportion correct (and standard deviations) at T1 and T2 (for directly encoded and non-encoded associations) for the TMR and non-TMR conditions are presented in Table 9.

To assess the effect of TMR on retention for open-loops, I performed a 2x2 (Session x TMR) ANOVA on retrieval accuracy at T1 and T2 (for open-loops not tested at T1). The ANOVA revealed a significant effect of Session, $F(1,28) = 26.79, p < .001, \eta_p^2 = .49$, with overall greater retrieval accuracy at T1 relative to T2. The ANOVA revealed no evidence for a significant effect of TMR, $F(1,28) = 1.15, p = .29, \eta_p^2 = .04$, nor interaction between Session and TMR, $F(1,28) = 1.47, p = 24, \eta_p^2 = .05$. Thus, in contrast to evidence for a positive effect of TMR on retention for single associations (see Cellini & Capuozzo, 2018 for review), I saw no evidence for TMR leading to greater retention for overlapping associations encoded as open-loops.

Table 9 Mean proportion correct (and standard deviations) at T1 (before sleep) and T2 (following sleep) for directly encoded pairs (tested and not tested at T1) and non-encoded pairs (associated with open-loops tested and not tested at T1) for the TMR and non-TMR conditions.

	T1	**T2**			
	Encoded	**Encoded**		**Non-encoded**	
Condition	**Tested at T1**	**Not tested at T1**	**Tested at T1**	**Not tested at T1**	**Tested at T1**
TMR	.62 (.19)	.54 (.20)	.64 (.23)	.45 (.25)	.51 (.28)
Non-TMR	.62 (.19)	.55 (.20)	.65 (.24)	.46 (.26)	.52 (.28)

A 2x2 ANOVA (Tested x TMR) for retrieval accuracy at T2 (for open-loops tested and not tested previously at T1) revealed a main effect of Tested, $F(1,28) = 47.12$, $p < .001$, $\eta_p^2 = .63$, with greater retrieval accuracy for open-loops previously tested at T1. No other significant effects were seen, Fs < 3.21, ps > .08.

3.4.2. Retrieval dependency

Mean proportion of joint retrieval (and standard deviations) for the data and independent model at T1 and T2 for the TMR and non-TMR conditions are presented in Table 10.

Direct comparisons between the proportion of joint retrieval in the data and independent model revealed no evidence for dependency in any of the experimental conditions, ts < 1.60, ps > .13. However, for retrieval at T1, open-loops in the non-TMR condition showed evidence for anti-dependency, $t(28) = 2.30$, $p < .01$, $d = .29$, $BF_{10} = 14.14$. Furthermore, for retrieval at T2 (for open-loops not tested previously at T1), there was borderline evidence for anti-dependency for open-loops in the TMR condition, $t(28) = 2.05$, $p = .05$, $d = 26$, $BF_{10} = 1.21$.

Table 10 Mean proportion of joint retrieval (and standard deviations) for the data and independent model at T1 (before sleep) and T2 (for open-loops tested and not tested previously at T1) (following sleep) for the TMR and non-TMR condition.

	T1		**T2**			
	Tested at T1		**Not tested at T1**		**Tested at T1**	
Condition	**Data**	**Independent**	**Data**	**Independent**	**Data**	**Independent**
TMR	.56 (.14)	.58 (.12)	.53 (.14)	.56 (.09)	.60 (.16)	.62 (.13)
Non-TMR	.52 (.15)	.56 (.13)	.57 (.13)	.57 (.08)	.62 (.18)	.62 (.15)

A 2x2 (Session x TMR) ANOVA on dependency (data – independent model) revealed no significant effects of Session, $F(1,28) = 1.77$, $p = .19$, $\eta_p^2 = .06$, nor TMR, $F(1,28) = .26$, $p = .61$, $\eta_p^2 < .01$, nor a significant interaction between Session and TMR, $F(1,28) = 2.9$, $p = .10$, $\eta_p^2 = .09$. Similarly, a 2x2 ANOVA (Tested x TMR) on dependency at T2 (for open-loops tested and not-tested previously at T1) revealed no significant main effects or interactions, *F*s < 2.09, *p*s > .15.

3.4.3. Inference accuracy

A 2x2 (Tested x TMR) ANOVA on inference performance revealed a significant main effect of Tested, $F(1,28) = 8.79$, $p < .01$, $\eta_p^2 = .24$, with greater inference for non-encoded pairs associated with open-loops previously tested at T1. No other significant effects were seen, *F*s < .74, *ps* > .39.

As I saw no evidence for TMR increasing inference, in addition to observing no evidence for increases in dependency as a function of TMR (or sleep more generally), the findings presented here suggest that increases in inference following sleep, as reported in Lau et al. (2010), are unlikely to be driven by mnemonic integration during sleep. The evidence presented here (and in Chapter 2) is perhaps more consistent with the proposal that successful inference is likely driven by ‘on-the-fly’ processes at the point of retrieval (Kumaran & McClelland, 2012), the probability of which, in this case, is increased due to greater retention for directly encoded associations previously tested at T1.

3.4.4. Retrieval accuracy for first and second encoded associations

So far, I have seen no evidence to suggest that TMR, relative to a non-TMR condition, leads to differences in retention for open-loops. However, in light of previous work showing that TMR during sleep may not benefit memory for overlapping associations equally (Antony et al., 2018; Oyarzún et al., 2017), I performed a 2x2x2 (Session x TMR x Encoding-order) ANOVA for retrieval accuracy at T1 and T2 (for open-loops not tested at T1) in order to assess whether the effect of TMR on retention may differ for open-loop associations depending on their encoding order. In other words, does TMR support or decrease retention depending on whether a given association was the first or second encoded association (e.g., *Barack-Obama-hammer,* then *kitchen-hammer*) of an open-loop?

Mean proportion correct (and standard deviations) for first and second encoded associations at T1 and T2 for the TMR and non-TMR conditions are presented in Table 11. Mean proportion correct at T1 and T2 (for associations not previously tested at T1) across TMR and encoding-order are presented in Figure 6.

The ANOVA revealed a significant effect of Encoding-order (first vs second encoded association), $F(1,28) = 35.57$, $p < .001$, $\eta_p^2 = .56$, with overall greater memory performance for the first, relative to the second encoded pair. Better memory performance for the first encoded, relative to the second encoded pair was seen at both T1, $t(28) = 5.92$, $p < .001$, $d = .75$, $BF_{10} > 1000$, and T2, $t(28) = 4.95$, $p < .001$, $d = 65$, $BF_{10} > 100$. This is consistent with previous work showing that the encoding of overlapping material is susceptible to proactive interference (see Postman, 1971 for review), such that the encoding (and/or retrieval) of the second association may be impaired by the presence of the first encoded association.

Table 11 Mean proportion correct (and standard deviations) for first and second encoded associations at T1 (before sleep) and T2 (after sleep) for the TMR and non-TMR condition.

	T1		**T2 (Not tested at T1)**		**T2 (Tested at T1)**	
Condition	**First**	**Second**	**First**	**Second**	**First**	**Second**
TMR	.70 (.21)	.54 (.21)	.64 (.21)	.44 (.23)	.69 (.24)	.58 (.25)
Non-TMR	.69 (.21)	.55 (.22)	.59 (.23)	.52 (.21)	.71 (.25)	.59 (.26)

The ANOVA also revealed a significant interaction between TMR and Encoding-order, $F(1,28) = 7.22$, $p = .01$, $\eta_p^2 = .21$. This interaction was further qualified by a three-way interaction between Session, TMR, and Encoding-order, $F(1,28) = 4.79$, $p = .04$, $\eta_p^2 = .15$. No other significant effects were seen, *F*s < 1.48, *p*s > .23. Further interrogation of the three-way interaction between Session, TMR, and Encoding-order revealed a significant interaction between TMR and Encoding-order at T2, $F(1,28) = 10.90$, $p < .01$, $\eta_p^2 = .28$, with better retrieval accuracy for the first encoded pair, $t(28) = 2.13$, $p = .04$, $d = .24$, $BF_{10} = 1.39$ (note though that this effect did not survive correction for multiple comparisons) and reduced retrieval accuracy for the second encoded pair, $t(28) = 3.53$, $p < .01$, $d = .37$, $BF_{10} = 24.08$, in the TMR vs non-TMR condition. No interaction between TMR and Encoding-order was seen at T1, $F(1,28) = .16$, $p = .70$, $\eta_p^2 = .01$. Thus, repeated memory reactivation during sleep supported retention for the first encoded associations and decreased retention for the second encoded associations. Note though that the former effect did not survive correction for multiple comparisons (and the Bayes factor only provided weak evidence in favour of this effect).

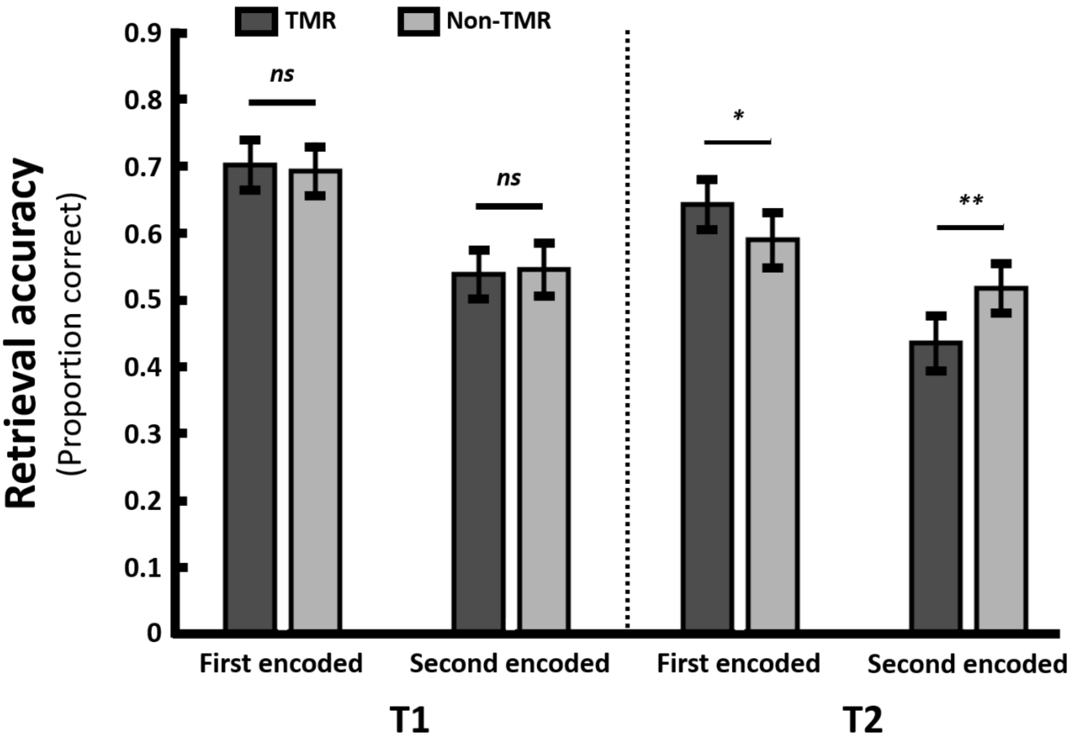

Figure 6 Mean proportion correct at T1 (before sleep) and T2 (following sleep) for the first and second encoded associations in the TMR and non-TMR condition. For T2, only trials where participants retrieved open-loops not previously tested at T1 are included. Error bars represent ± 1 standard error. *ns* = not significant. ** $p < .01$, * $p < .05$ (uncorrected).

To assess whether the effect of TMR differed for open-loops tested and not-tested at T1, the analysis above was repeated using a 2x2x2 ANOVA (Tested x TMR x encoding-order) but focussing exclusively on retrieval at T2 for pairs previously tested and not-tested at T1. This analysis revealed a significant main effect of Tested, $F(1,28) = 47.11$, $p < .001$, $\eta_p^2 = .63$, with overall greater accuracy for open-loops tested previously at T1. The ANOVA again revealed a significant effect of Encoding-order, $F(1,28) = 27.11$, $p < .001$, $\eta_p^2 = .49$, with greater retrieval accuracy for the first, relative to second encoded associations. Further, the ANOVA revealed a significant three-way interaction between Tested, TMR, and Encoding-order, $F(1,28) = 6.11$, $p = .02$, $\eta_p^2 = .18$. No other significant effects were observed, Fs < 3.21, ps > .08. However, in contrast to associations not tested previously at T1 (see results above), no significant interaction between TMR and Encoding-order was seen at T2 for associations previously tested at T1, $F(1,28) = .03$, $p = .87$, $\eta_p^2 < .01$. Thus, prior testing at T1 appeared to reduce the influence of TMR during sleep, relative to associations not tested prior to sleep.

3.5. General discussion

In Chapter 2, I provided evidence that sleep, relative to wakefulness, increases inference for non-encoded, but overlapping associations. This is consistent with evidence for increases in inference following sleep (Lau et al., 2010). However, in contrast to the proposal that inference is driven by the formation of a generalised representation during sleep (Lau et al., 2010), the evidence presented in Chapter 2 suggested that inference is perhaps more likely to be supported by ‘on-the-fly’ processes at retrieval (Banino et al., 2016), as no evidence for mnemonic integration (as measured by retrieval dependency) was seen follow sleep. Here I used TMR to assess the role of memory reactivation during sleep, rather than sleep more generally (as in Chapter 2), in supporting inference and integration.

In contrast to the proposal that memory reactivation during sleep strengthens the association between overlapping information (Diekelmann & Born, 2010), I found no evidence to suggest that TMR during sleep promotes integration of open-loops (as measured by dependency), nor does it appear to improve participants’ ability to make inferences between related information (i.e., inferring that B-C are related after encoding A-B and A-C pairs). Instead, I provided evidence that TMR decreases memory performance for the second encoded association, and supported memory for the first encoded association of an open-loop (but this latter effect did not survive correction for multiple comparisons, and the Bayes factor only provide weak evidence in favour of this effect). Interestingly, the effect of TMR was only seen for pairwise associations that had not been tested prior to sleep.

Despite the long-standing proposal that sleep reduces forgetting (Jenkins & Dallenbach, 1924), recent work has shown that memory reactivation during sleep can lead to forgetting of overlapping (or competing) memories, depending on the interval between encoding (Oyarzún et al., 2017) or their expected future reward (Antony et al., 2018). In line with this, the evidence presented here showed that memory reactivation during sleep can lead to forgetting of overlapping associations depending on their encoding order, even when the interval between the encoding of two pairwise associations is relatively short as compared to that used in Oyarzún et al. (2017).

At immediate test, retrieval accuracy for the first encoded associations was greater than that for the second encoded associations. If higher retrieval accuracy relates to greater associative strength, then it is possible that the first and second encoded associations differ in their mnemonic strength, presumably as an effect of proactive interference at encoding (see Postman, 1971 for review). Interestingly, previous work has suggested that sleep reduces proactive interference (Abel & Bäuml, 2014). There is also evidence to suggest that memory consolidation during sleep is sensitive to differences at encoding (Schmidt et al., 2006), and may prioritise the consolidation of weakly encoded information (Cairney et al., 2016; Drosopoulos et al., 2007; Schapiro et al., 2018). However, I provided evidence that retention for the second (or weakly) encoded associations decreased at a greater rate in the TMR, relative to the non-TMR condition. Although the precise mechanisms that underlie this pattern of forgetting is unclear, it is possible that re-exposing participants to specific auditory stimuli during sleep is more likely to lead to the reactivation of the first encoded association (given that greater memory performance is seen for this association prior to sleep) and results in the forgetting of the second encoded association of an open-loop. This would be consistent with evidence showing that the repeated retrieval of a ‘target’ memory during wakefulness can lead to retrieval-induced forgetting of a ‘competing’ memory representation (Anderson et al., 1994; Wimber et al., 2015).

Although evidence for better retention for the first encoded pairs was seen in the TMR, relative to non-TMR condition, this effect did not survive correction for multiple comparisons. It is possible that any retention benefits may become more apparent over time as the effects of TMR do not always arise immediately following sleep (Cairney et al., 2018). However, retrieval-induced forgetting is thought to be due to the weakening of a competing memory, rather than the strengthening of the target memory (Anderson, Bjork, & Bjork, 2000). If TMR is more likely to lead to the reactivation of the first (or strongly) encoded association, then it is possible that the findings observed here may reflect a similar retrieval-induced forgetting effect that leads to decreases in retention for the second (or weakly) encoded association of an open-loop.

Retrieval-induced forgetting is associated with activity in the prefrontal cortex (Kuhl et al., 2007; Kuhl, Kahn, Dudukovic, & Wagner, 2008; Wimber et al., 2015, 2008), suggesting that retrieval-induced forgetting may be induced by a top-down, inhibitory mechanism (Anderson, 2003). Consistent with Oyarzún et al. (2017), the evidence presented here suggests that forgetting of a memory representation following repeated memory reactivation can occur during sleep, presumably in the absence of any conscious executive control. According to computational accounts (Norman, Newman, Detre, & Polyn, 2006; Norman, Newman, & Detre, 2007), retrieval-induced forgetting is due to differences in activation at retrieval, such that forgetting occurs when levels of activation differ between a ‘target’ and ‘competitor’ memory. This computational account of retrieval-induced forgetting may possibly explain the findings presented here, under the assumption that TMR is more likely to lead to greater levels of activation for the first, relative to second encoded association. Further work is needed to assess this possibility. However, if the findings reported here are due to differences in the level of activation for overlapping associations during TMR, then it is possible that the effect of TMR may differ when the mnemonic strength of associations from the same ‘event’ are relatively similar, as appears to be the case when overlapping pairwise associations are encoded in a closed-loop (Horner & Burgess, 2013, 2014).

It has been proposed that systems consolidation during sleep is supported by a set of hierarchically nested oscillations (Buzsáki, 1996). Slow-oscillations (~1Hz) originating in the neocortex are thought to provide the temporal frame for memory reactivation during sleep (Sirota & Buzsáki, 2005). Slow-oscillations extend into to the thalamus (Massimini et al., 2004) where they are proposed to trigger the onset of sleep spindles (~12-15Hz) (Mölle et al., 2002). Importantly, increased densities of slow oscillations (Marshall et al., 2006) and sleep spindles (Bergmann et al., 2012; Mednick et al., 2013) have been shown to support memory consolidation during sleep. It is possible that the findings presented here are characterised by differences in the pattern of slow-oscillations and/or sleep spindles following auditory stimulation during sleep.

Oyarzún et al. (2017) showed that increased oscillatory activity in the beta band (15-30Hz) following TMR presentation was associated with forgetting. Consistent with this, increases in beta activity have been shown to predict retrieval-induced forgetting in wakefulness (Waldhauser, Johansson, & Hanslmayr, 2012). In contrast, theta activity has been associated with successful memory retrieval in wake (Düzel, Neufang, & Heinze, 2005) and to predict memory retention for associations between native words and their foreign translations following TMR (Schreiner, Göldi, & Rasch, 2015; Schreiner & Rasch, 2015). However, theta power has also been associated with retrieval-induced forgetting for overlapping memories in wakefulness. (Hanslmayr, Staudigl, Aslan, & Bäuml, 2010). As such, it is possible that differences in theta and beta activity following the offset of auditory stimulation during TMR may predict whether weakly and/or strongly encoded associations are retained or forgotten. Further work is needed to assess whether distinct oscillatory responses during TMR are associated with differences in retention for overlapping memory representations following sleep.

Interestingly, an effect of TMR on retention was not observed for open-loops tested prior to sleep. Retrieval practice has been shown to have robust effects on long-term retention (see

Roediger & Butler, 2011 for review). One proposal (Kornell, Bjork, & Garcia, 2011) holds that prior testing increases the strength of a memory representation, such that the memory may stay above a retrieval threshold after both a short and prolonged retention interval. It is possible that testing prior to sleep increases the mnemonic strength of open-loops, such that even weakly encoded associations of an open-loop become resistant to the effect of TMR during sleep. This would be consistent with evidence showing that sleep promotes the retention of memories subject to restudy, but not those retrieved prior to sleep (Bäuml, Holterman, & Abel, 2014), suggesting that memory representations retrieved prior to sleep may be retained at a sufficiently high level to mitigate the role of sleep in memory consolidation. Curiously, most of the previous work assessing the effect of TMR during sleep on long-term memory has tended to focus on the retention of information tested prior to sleep. It is possible that the effects reported in previous work are obscured by the effect of prior testing.

3.6. Conclusion

I provided evidence that memory reactivation during sleep (as a function of TMR) does not support memory for overlapping associations equally. TMR during sleep was shown to lead to decreases in retention for the second (or weakly) encoded association of an open-loop; possibly as an effect of proactive interference at encoding, but this was dependent on whether an association was tested prior to sleep or not. These findings are consistent with previous work showing that TMR can, in some instances, lead to forgetting (Antony et al., 2018; Oyarzun et al., 2017).

3.7. Acknowledgments

To ensure greater sleep scoring accuracy, the sleep data was scored separately by the author and Dr Marcus Harrington. I wish to thank Marcus for his assistance with sleep scoring.

4. Holistic reinstatement of closed-loops following a period of forgetting and consolidation

4.1. Abstract

Research has shown that episodic memory retrieval is associated with the holistic, neocortical reinstatement of all event elements; a process supported by the hippocampus. However, it is not clear whether event-based representations continue to be retrieved in a holistic manner, nor whether the hippocampus continues to support holistic retrieval following a period of forgetting and consolidation. Here, I used fMRI to assess whether the retrieval of closed-loops is associated with the reinstatement of all event elements, even those incidental to the retrieval task, following a 24-hour delay. I provided evidence that closed-loops continue to be reinstated holistically in the neocortex following a period of forgetting and consolidation. However, while the strength of reinstatement for incidental closed-loop elements retrieved immediately after encoding correlated with activity in the hippocampus, the degree of reinstatement for elements incidental to retrieval did not correlate significantly with activity in the hippocampus for closed-loops retrieved 24-hours after encoding. These findings provide evidence that event-based representations continue to be retrieved in a holistic manner but suggest that the role of the hippocampus in holistic retrieval may decrease over a period of forgetting and consolidation.

4.2. Introduction

The process of retrieving an episodic memory has long been thought to involved the holistic retrieval of all aspects of a past event (Tulving, 1983). Despite this long-standing proposal, only recently has evidence emerged to demonstrate that event-based representations are reinstated in their entirety during retrieval (Horner et al., 2015). Although evidence for neocortical reinstatement of associative information has been shown previously (Johnson & Rugg, 2007; Polyn et al., 2005; Wheeler et al., 2000; Woodruff et al., 2005), Horner et al. (2015) showed that after learning the associations between, for instance, *kitchen-hammer*, *kitchen-Barack Obama*, and *Barack Obama-hammer*, cuing participants with *kitchen* and asking them to retrieve *Barack Obama* also led to increased activity in neocortical regions associated with the remaining ‘non-target’ element *hammer*, despite no explicit requirements to retrieve this element. The holistic reinstatement of all event elements in the neocortex - irrespective of whether the element is require to be retrieved or not - is consistent with Tulving's (1983) proposal that episodic memory is associated with holistic retrieval, and provides evidence for the idea that recollection is associated with all-or-none retrieval (Norman & O’Reilly, 2003)

Critically, Horner et al. (2015) showed that the strength of reinstatement for ‘non-target’ elements incidental to retrieval was correlated with activity in the hippocampus, providing evidence for all-or-none retrieval occurring via the computational process of pattern completion (see also Grande et al., 2019). This is consistent with the proposed role of pattern completion in models of the hippocampus as an associative memory network (Gardner-Medwin, 1976; Hopfield, 1982; McClelland et al., 1995; Treves & Rolls, 1992). However, standard theories of consolidation (e.g., McClelland et al., 1995; Squire & Alvarez, 1995) have proposed that event-based representations can be retrieved independently of the hippocampus after a period of consolidation. This proposal is motivated by observations that although patients with hippocampal damage present with amnesia (Scoville & Milner, 1957; Vargha-Khadem, 1997), they often show a temporal gradient to their impairment, with remote events being relatively spared compared to more recent events (Schnider, Bassetti, Gutbrod, & Ozdoba, 1995; but see Sanders & Warrington, 1971). Consistent with standard theories of consolidation (e.g., McClelland et al., 1995; Squire & Alvarez, 1995), hippocampal activity associated with successful retrieval has been shown to decrease over a one-day interval (Takashima et al., 2006). Given this finding, it is possible that the correlation between activity in the hippocampus and the strength of the neocortical reinstatement effect observed in Horner et al. (2015) may decrease as a function of time. Consistent with this, connectivity between the hippocampus and neocortical regions associated with successful retrievals of object-location pairs has been shown to decrease over a 24-hour interval (Takashima et al., 2009)

In Horner et al. (2015), neocortical activity associated with the reinstatement of all event elements, even those incidental to retrieval, was only seen for ‘events’ where all possible pairwise associations had been learnt in a closed-loop associative structure. This is consistent with previous behavioural findings showing that while the separate encoding of three overlapping pairwise associations in a closed-loop results in retrieval dependency for all elements of an event, dependency is not seen when only two out of the three possible associations (e.g., *kitchen* and *hammer*, *kitchen* and *Barack Obama*, but not *Barack Obama* and *hammer*) are encoded in an open-loop (Horner & Burgess, 2014). Thus, it is only in the closed-loop condition where the pairwise associations appear to be bound into a coherent event representation (Horner et al., 2015; Horner & Burgess, 2014). Consistent with this, dependency for closed-loops is similar to that seen for three-element events formed on a single encoding trial (Horner & Burgess, 2014), suggesting that the closed-loop associative structure is critical for the formation of a coherent event-based representation in the absence of a distinct spatial-temporal context; as typically attributed to episodic memory (Tulving, 1983).

In Chapter 2, I used the presence of dependency to infer the coherence of an underlying event-based representation and showed that dependency is retained after a period of forgetting. If participants learn the associations between *Barack Obama*, *kitchen*, and *hammer* in a closed-loop (as in Horner et al., 2015; Horner & Burgess, 2014), the elements that make up this event do not appear to be forgotten independently, but instead tend to be forgotten in an all-or-none manner. This pattern of forgetting is consistent with behavioural evidence showing that recollection – a process thought to be supported by the hippocampus (Diana et al., 2010; Mayes et al., 2007) - decreases as a function of decay, rather than interference, between encoding and retrieval (Sadeh et al., 2016), and provides evidence that the rate of decay is relatively uniform within an 'event'.

Here I asked, as event-based representations are forgotten, do we see decreases in neocortical activity associated with the reinstatement of elements incidental to retrieval. If event-based representations, or closed-loops more specifically, retain their coherence over a period of forgetting (as shown in Chapter 2) then evidence for neocortical reinstatement for elements incidental to retrieval should be seen not only immediately after encoding but also following a delay. For instance, if a participant learns the associations between *Barack Obama*, *kitchen*, and *hammer* and this closed-loop retains its coherence following a period of forgetting and consolidation, cueing the participant with *kitchen* should lead to the reinstatement of both *Barack Obama* and *hammer*, regardless of which of the elements are required to be retrieved. Note that the argument here is that forgetting within an event should be sufficiently uniform for event-based representations to retain their coherence at retrieval. It is entirely possible that event-based representations change in some continuous manner over a period of forgetting. However, if the rate of forgetting within an event is sufficiently uniform for all elements within the same event to fall either above or below a given retrieval threshold (e.g., Norman & O'Reilly, 2003; Yonelinas, 1994) then closed-loops should continue to be retrieved in a holistic manner following a period of forgetting.

I also asked whether the correlation between activity in the hippocampus and the reinstatement of elements incidental to retrieval (as observed in Horner et al. (2015)) differs depending on whether closed-loops are retrieved immediately after encoding or after a 24-hour delay (if closed-loops indeed continue to be reinstated holistically). This is because although some evidence has shown that the role of the hippocampus in retrieval decreases over time (Niki & Luo, 2002; Piefke et al., 2003), a growing number of studies have supported a persistent role for the hippocampus in the recall of rich contextual 'events', irrespective of the age of the memory (Addis et al., 2004; Bonnici et al., 2012; Gilboa et al., 2004; Maguire et al., 2001; Ryan et al., 2001). This is consistent with the proposal that while event-based representations can gradually become represented in the neocortex in a gist-like manner, the retrieval of event-based representations, at least in their rich contextual form, will continue to be dependent on the hippocampus, regardless of the interval between encoding and retrieval (Nadel & Moscovitch, 1997; Winocur & Moscovitch, 2011).

I used a similar design to that developed in Horner and Burgess (2014) (see Figure 7). Across two encoding sessions, separated by ~24 hours, participants learnt a series of multi-element triplets; each consisting of three elements (famous people, locations and common objects). Triplets were 'built up' over two/three separate encoding trials with each trial consisting of the presentation of one of three possible pairwise associations for a triplet. This allowed for the formation of triplets with different associative structures: 'closed-loops' where all the pairwise associations are encoded, or 'open-loops' where only two out of the three possible pairwise associations are encoded. Although closed- and open-loops differ in the number of pairwise associations, previous findings have shown that the presence or absence of holistic retrieval during cued-recognition depends on whether pairwise associations form part of a closed- or open-loop, even when closed- and open-loops are matched in terms of pairwise associations, but not in the number of elements (Horner et al., 2015; Horner & Burgess, 2014). However, the critical question here is whether holistic reinstatement is seen for closed-loops retrieved immediately after encoding and following a 24-hour delay. As such, open-loops serve as a baseline condition to assess the presence or absence of holistic retrieval for closed-loops retrieved immediately after encoding and following a period of forgetting and consolidation.

At retrieval, the learnt associations were tested in both directions (e.g., cue: *kitchen*, target: *Barack Obama*; cue: *Barack Obama*, target: *kitchen*) using six-alternative forced choice among elements of the same category from different closed- and open-loops. FMRI was used to assess neocortical activity for closed- and open-loops retrieved immediately after encoding and following a 24-hour delay. As in Horner et al. (2015), the analyses focused on identifying differences in retrieval activity between closed- and open-loops for elements required to be retrieved (i.e., cues and targets) and those incidental to retrieval (i.e., non-targets). As activity in the hippocampus has been shown to correlate with the degree of incidental reinstatement when closed-loops are retrieved immediately after encoding (Horner et al., 2015), I also assessed whether activity in the hippocampus correlated with the neocortical reinstatement of elements incidental to retrieval for closed-loops immediately after encoding and following a 24-hour delay. In short, the experiment asked (1) whether closed-loops are associated with the reinstatement of event elements incidental to retrieval after a period of forgetting and consolidation, and (2) whether the hippocampus correlates with the degree of incidental reinstatement for closed-loops retrieved following a period of forgetting and consolidation.

4.3. Methods

4.3.1. Participants

An effect size $d = .67$ for detecting a significant non-target reinstatement effect was calculated from previous work (Horner et al., 2015), with $n = 26$. Cohen’s d was calculated from the reported t-statistic divided by the square root of the within-subject sample size (Lenth, 2006).

However, as forgetting is expected to occur when retrieval follows a delay, retrieval accuracy at delay is expected to be lower than in Horner et al. (2015). Pilot data, with *n* = 12, revealed ~14% decrease in retrieval accuracy for closed- and open-loops retrieved following a 24-hour delay, relative to performance (.64) seen in Horner et al. (2015), leading to an adjusted effect size *d* = .62. Using G*Power (Faul et al., 2009) and an effect size *d* = .62, I estimated that a sample size of *n* = 30 was required to detect a significant non-target reinstatement effect, if one is present, at a power of .90 and $\alpha = .05$.

Thirty-three, right-handed participants were recruited from the University of York student population. All participants gave written informed consent and took part in exchange for course credit or monetary compensation. Participants had either normal to corrected-to-normal vision and reported no history of neurological or psychiatric illness. Data from three participants could not be included in the final sample due to an excess of motion related artefacts in the imaging data (*n* = 1) and memory accuracy not greater than chance (< ~16.7%) (*n* = 2). Accordingly, the analyses included 30 participants (23 female) with a mean age of 22.23 years ± 4.49 *SD*.

The experiment was approved by the York Neuroimaging Centre (YNiC) ethics committee, University of York, UK.

4.3.2. Materials

The stimuli consisted of 72 locations (e.g., *Kitchen*), 72 famous people (e.g., *Barack Obama*), and 72 common objects (e.g., *hammer*). From these, 72 random location-person-object triplets were generated for each participant. The term ‘triplet’ is used to refer to the three elements (location, person, object) that were assigned to the same associative structure (closed- or open-loop). 36 triplets were randomly assigned to the within-subject experimental conditions closed- vs open-loop. From these, 18 open- and closed-loops were randomly assigned to the experimental conditions delay vs no-delay. Thus, 18 triplets were randomly assigned to each of the four within-subject experimental conditions: (1) open-loop, delay; (2) closed-loop, delay; (3) open-loop, no-delay; and (4) closed-loop, no-delay.

For closed-loops, all three possible pairwise associations for each triplet were encoded. For open-loops, only two out of the three possible pairwise associations were encoded. Triplets were never presented together at encoding or retrieval; only specific pairwise associations were encoded and retrieved for each triplet, depending on whether they were open- or closed-loops. Note that open-loops are equated to the closed-loops in the number of elements, but not in the number of associations. It has previously been shown (Horner et al., 2015; Horner & Burgess, 2014) that dependency is not seen for open-loops when three overlapping associations are encoded in an associative chain (e.g., *kitchen-hammer*, *kitchen-Barack Obama*, *Barack Obama-dog*), controlling for the number of associations (but not the number of elements). Any differences in retrieval between the two conditions in the current experiment are therefore unlikely to be driven by differences in the number of associations.

4.3.3. Procedure

The experiment consisted of two encoding sessions and a single retrieval session. Both encoding sessions took place outside the scanner and were separated by ~24 hours (*M* = 23 hr, 36 min), including overnight sleep. The retrieval session took place inside the scanner immediately after the second encoding session.

Encoding. During encoding, participants were presented with specific pairwise associations for each of the 72 triplets. Participants learnt one pairwise association per encoding trial. All pairwise associations were presented on a computer screen as words, with one item to the left and one item to the right of fixation. The left/right assignment was randomly chosen on each trial. The words remained on screen for 6s. Participants were instructed to imagine, as vividly as possible, the items interacting in a meaningful way for the full 6s. Each word-pair presentation was preceded by a 500-ms fixation and followed by a 500-ms blank screen. For

36 out of the 72 triplets, participants encoded all three possible pairwise associations; forming closed-loops. For the remaining 36 triplets, participants encoded two out of three possible pairwise associations; forming open-loops. 18 closed- and 18 open-loops were randomly assigned to one of the two encoding sessions, making a total of 90 encoding trials per encoding session.

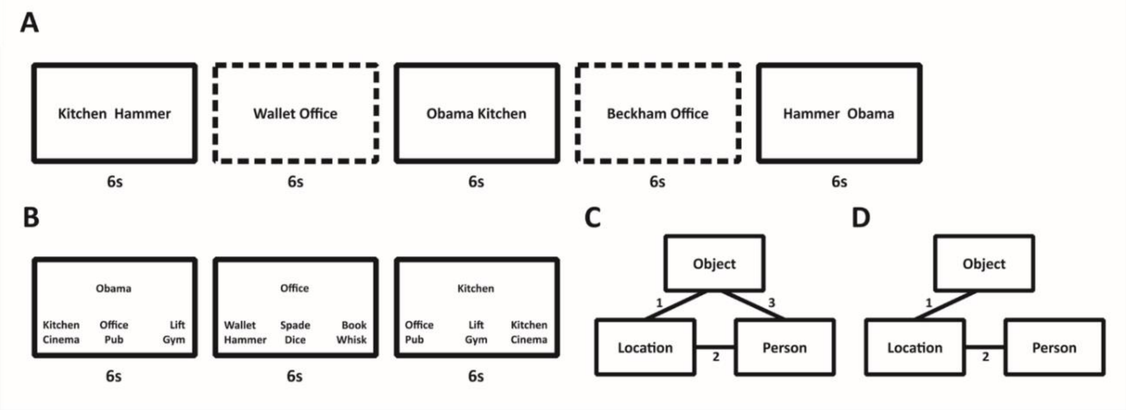

Figure 7 Experimental design. **(A)** Encoding. Participants saw multiple pairwise associations. They imagined each association 'interacting in a meaningful way as vividly as possible' for 6s. Each association was preceded by a 500-ms fixation and followed by a 500-ms blank screen. Participants encoded open- and closed-loop pairwise associations in an intermixed manner. In the open-loop condition, participants did not encode the third and final association (e.g., *wallet-Beckham*) (see **D**). Solid and dotted lines were not presented but highlight closed- (solid lines) and open-loops (dotted lines). **(B)** Retrieval. Participants were presented with a single cue and required to retrieve one of the other elements from the same triplet from among five foils (elements of the same type from other triplets) in 6s. Each cued-recognition trial was preceded by a 1s fixation and followed by a 1s blank screen. **(C)** The associative structures of closed-loops with example encoding order for three pairwise associations (numbers 1-3). **(D)** The associative structure of open-loops with example encoding order for the two pairwise associations (numbers 1-2). The third and final association (i.e., person-object in this example) is not shown to the participants.

Each encoding session consisted of three blocks of 18, 36, and 36 trials, respectively, with one pair from each triplet being presented during each block (participants were not made aware of this structure). During the first encoding block, only pairwise associations for closed-loops were presented. This ensured that the duration between encoding the last pairwise associations and retrieval was consistent for closed- and open-loops from each encoding session. In blocks two and three, the open- and closed-loop associations were randomly presented in an intermixed manner. The presentation order was randomised for each block. A break of 10s would follow every 18 encoding trials.

Each open-loop consisted of a common element (e.g., if participants learnt location-person, and location-object, then location would be the common element). For each encoding session, six out of the total 18 open-loops were pseudo-randomly assigned to one of the three possible common elements (i.e., locations, people, or objects). The presentation order for open-loops across encoding blocks two and three was: (1) person-location, location-object; (2) location-object, object-person; (3) object-person, person-location. For each encoding session, closed-loops were pseudo-randomly rotated in the same manner. The presentation order for closed-loops across the three encoding blocks was: (1) person-location, location-object, object-person; (2) location-object, object-person, person-location; (3) object-person, person-location, location-object.

Retrieval. During retrieval, participants performed a six-alternative forced-choice cued-recognition task. On each trial, a cue and six possible targets were presented simultaneously on screen. The cue was presented in the middle of the screen with six possible targets; one target and five foils from the same category (e.g., if the target was *hammer*, the five foils would be other randomly selected objects from the other triplets, regardless of closed- vs open-loop or delay vs no-delay status), presented in two rows of three below the cue. Participants had a maximum of 6s to respond with a button response, that corresponded to the targets position on screen, and were instructed to be as accurate as possible in the time given. The position of the correct target was randomly selected on each retrieval trial. The cue and six targets were presented until a response was made (up to a maximum of 6s). Missing responses (i.e., responses that fell outside the 6s response window) (M = 2.52% ± 2.26 SD) were treated as incorrect trials.

Each closed- and open-loop was tested with one of the cue-target associations (e.g., cue: person, target: location) in both directions across two retrieval blocks (total retrieval trials =

360). The presentation order was optimised to measure univariate BOLD activity in each of the four within-subject experimental conditions (i.e., open- vs closed-loop, delay vs no-delay) (optimisation algorithm available at https://osf.io/eh78w/). To avoid adaptation effects, cue-target associations from the same closed- or open-loop were never presented on successive trials and each block contained 18 null trials that each lasted 6s. In each block, cue-target associations belonging to nine closed- and nine open-loops encoded during the first encoding session, and nine closed- and nine open-loops encoded immediately before retrieval were tested, making a total of 180 retrieval trials, in addition to 18 null trials, per block. Each trial was preceded by a 1s fixation and followed by a 1s blank screen.

4.3.4. Behavioural analyses

A 2x2 (Loop x Delay) ANOVA was used for the behavioural analysis of retrieval accuracy and dependency (see Chapter 2 for description of how retrieval dependency is modelled), with the within-subject factor Loop referring to whether the triplets form closed- or open-loops, and Delay referring to whether closed- or open-loops were encoded immediately or 24 hours prior to retrieval. Note that for the dependency analysis, the dependent variable refers to the difference in the proportion of joint retrieval in the data and independent model (i.e., dependency). For the analysis of retrieval dependency, I also report *t*-tests comparing the proportion of joint retrieval in the data against the independent model.

Alpha was set to .05 (two-tailed) for all statistical tests. For each ANOVA, a η_p^2 effect size is reported, and for *t*-tests I report a Cohen's *d* effect size reflecting the mean difference between conditions divided by the pooled standard error (Lakens, 2013). All analyses (including fMRI; see below) were conducted using MATLAB (R2019a).

4.3.5. fMRI acquisition

All functional and structural volumes were acquired on a 3-Tesla Siemens MAGNETOM Prisma scanner equipped with a 64-channel phase array head coil. T2*-weighted slices were acquired with echo-planar imaging (EPI). 48 axial slices (~0° tilt to the AC-PC line) per volume were acquired in an interleaved order with the following parameters: acquisition matrix = 64 x 64, repetition time = 1200-ms, echo time = 26-ms, flip angle = 75°, slice thickness = 3-mm, in-plane resolution = 3 x 3-mm, multi-band acceleration factor = 2. To allow for T1 equilibrium, the first three EPI volumes were acquired prior to the task and then discarded. As the retrieval phase varied in length across participants given each cued-recognition trial was self-paced (up to a maximum of 6s), the number of acquired volumes differed across participants. Note that the two retrieval blocks corresponded to two separate functional scanning runs. The mean number of volumes acquired during retrieval was 788.07 (range 614-955) and 760.70 (range 615-941) for the first and second retrieval block, respectively. A field map was acquired to allow for the correction of geometric distortions due to field inhomogeneities. For the purposes of co-registration and image normalisation, a whole-brain T1-weighted structural scan was acquired with a 1-mm^3 resolution using a magnetisation-prepared rapid gradient echo (MPRAGE) pulse sequence.

4.3.6. fMRI analyses

Preprocessing. Image pre-processing was performed in SPM12 (https://www.fil.ion.ucl.ac.uk/spm/). EPI images were corrected for field inhomogeneity based on geometric distortions, and spatially realigned to the first image in the time series. EPI images were spatially normalised to MNI space with transformation parameters derived from warping each participant’s structural image to a T1-weighted average template image (using the DARTEL toolbox; Ashburner, 2007). EPI images were spatially smoothed with an isotropic 8-mm FWHM Gaussian kernel prior to analysis.

General analysis. At the first level, BOLD activity was modelled by a boxcar function from cue/target onset to the time a response was made (up to a maximum of 6s). The predicted BOLD response was then convolved with a canonical hemodynamic response function to produce regressors of interest. In addition to the main regressors of interest, all first-level models included 13 regressors representing the movement parameter, and their derivatives, estimated during realignment, and a volume exclusion regressor. The volume exclusion regressor was defined as any volumes where the movement derivatives exceeded 1.5-mm translation (i.e., 0.5 x slice thickness) or 1° rotation. Parameter estimates for each regressor of interest were included in the second level analyses to identify consistent effects across participants. Unless otherwise stated, all second level models explicitly modelled subject effects. All effects reported outside the hippocampus are $p < .05$ FWE whole-brain corrected.

Neocortical reinstatement. The first level model included 24 regressors that related to all cue-target pairs (i.e., people-locations, people-objects, locations-people, locations-objects, objects-people, and objects-locations) for each of the four experimental conditions (open vs closed and delay vs no-delay). At the second level, regions that showed greater BOLD response to cueing/retrieving each element-type (i.e., cue/target), relative to when an element-type was not cued or required to be retrieved (i.e., non-target), collapsed across the four experimental conditions, were identified. This revealed three separate regions of interest across the three element-types. For each region of interest – defined as a 9-mm radius sphere centred on the peak coordinate - BOLD responses for each individual across all 24 regressors were extracted. The BOLD response when the element associated with a given region was either the cue, target, and non-target across the four experimental conditions was then calculated. As such, estimates of the BOLD response across 12 conditions (i.e., cue, target and non-target x closed- and open-loop x delay and no-delay) were obtained for each of the three regions of interest.

Non-target reinstatement and correlated activity. The first level model included four regressors corresponding to trials from each of the four experimental conditions (i.e., open-loop, delay; closed-loop, delay; open-loop, no-delay; and closed-loop, no-delay). At the second level, two separate models were created, one for the delay and one for the no-delay condition. Each model included a single regressor corresponding to the contrast between closed- and open-loops (i.e., open-loop vs closed-loops for the respective delay or no-delay condition), and a further covariate that corresponded to the difference in non-target activity between the closed- and open-loop condition (for the respective delay or no-delay condition). The covariate was calculated by taking the mean difference between the closed- and open-loop conditions for the non-target element across the three regions of interest. This was done separately for the delay and no-delay conditions. Given the *a priori* hypothesis regarding the hippocampus, statistical effects in the bilateral hippocampus are $p < .05$ FWE SVC within a bilateral hippocampus mask. The bilateral hippocampus mask for SVC was created using the WFU PickAtlas toolbox (Maldjian, Laurienti, Kraft, & Burdette, 2003) with the hippocampus defined from the Automated Anatomical Labelling atlas (Tzourio-Mazoyer et al., 2002). Given that this analysis focused on identifying differences across participants, no subject effects were included.

4.4. Results

4.4.1. Behaviour

Retrieval accuracy. Mean proportion correct (and standard deviations) across the four experimental conditions are presented in Table 12.

A 2x2 (Loop x Delay) ANOVA revealed significant effects of Loop, $F(1,29) = 21.91$, $p < .001$, $\eta_p^2 = .43$, and Delay, $F(1,29) = 18.22$, $p < .001$, $\eta_p^2 = .39$, with greater overall accuracy for closed, relative to open-loops, and for closed- and open-loops retrieved immediately after encoding, relative to those retrieved following a 24-hour delay. Consistent with findings reported in Chapter 2, the ANOVA also revealed a significant interaction between Loop and Delay, $F(1,29) = 6.48$, $p = .02$, $\eta_p^2 = .18$, with a greater difference in retrieval accuracy between closed- and open-loops at delay, relative to no-delay.

Table 12 Mean proportion correct (and standard deviations) and mean proportion of joint retrieval (and standard deviations) for the data and independent model at no delay (i.e., encoded immediately prior to retrieval) and delay (i.e., encoded 24 hours prior to retrieval) for closed- and open-loops.

		Proportion correct	**Proportion of joint retrieval**	
Delay	**Loop**		**Data**	**Independent model**
No delay	Open-loop	.61 (.21)	.62 (.14)	.59 (.13)
	Closed-loop	.67 (.24)	.70 (.15)	.67 (.15)
Delay	Open-loop	.52 (.21)	.59 (.14)	.58 (.13)
	Closed-loop	.62 (.25)	.69 (.13)	.64 (.14)

Retrieval Dependency. The mean proportion of joint retrieval (and standard deviations) for the data and independent model for closed- and open-loops at delay and no-delay are presented in Table 12. Figure 8 shows mean dependency (i.e., data – independent model) across delay vs no-delay and closed- vs open-loop.

Evidence for dependency was seen for closed-loops at both delay, $t(29) = 5.30$, $p < .001$, $d = .38$, and no-delay, $t(29) = 3.22$, $p = < .01$, $d = 21$. No evidence for dependency was seen for open-loops at either delay, $t(29) = 1.51$, $p = .14$, $d = .10$, or no-delay, $t(29) = 1.86$, $p = .07$, $d = .19$ (but note the marginal *p*-value for open-loops at no-delay). A 2x2 (Loop x Delay) ANOVA revealed no significant effects of Loop nor Delay, *F*s < 2.98, *ps* > .09, nor a significant interaction between Loop and Delay, $F(1,29) = 35.2$, $p = .07$, $\eta_p^2= 0.11$ (but note the marginal *p*-value). This interaction, albeit non-significant, appeared to be driven by a lack of a difference in dependency between closed- and open-loops at no-delay, $t(29) = .27$, $p = .79$, $d = .07$. Significantly greater dependency for closed, relative to open-loops, was seen at delay, $t(29) = 3.06$, $p < .01$, $d = .75$. It is possible that participants have inferred the associative nature of the overlapping pairs following the first encoding session (i.e., delay), leading to increases in dependency for open-loops encoded during the second encoding session (i.e., no-delay). However, it is important to note that dependency is seen for closed-loops at both delay and no-delay.

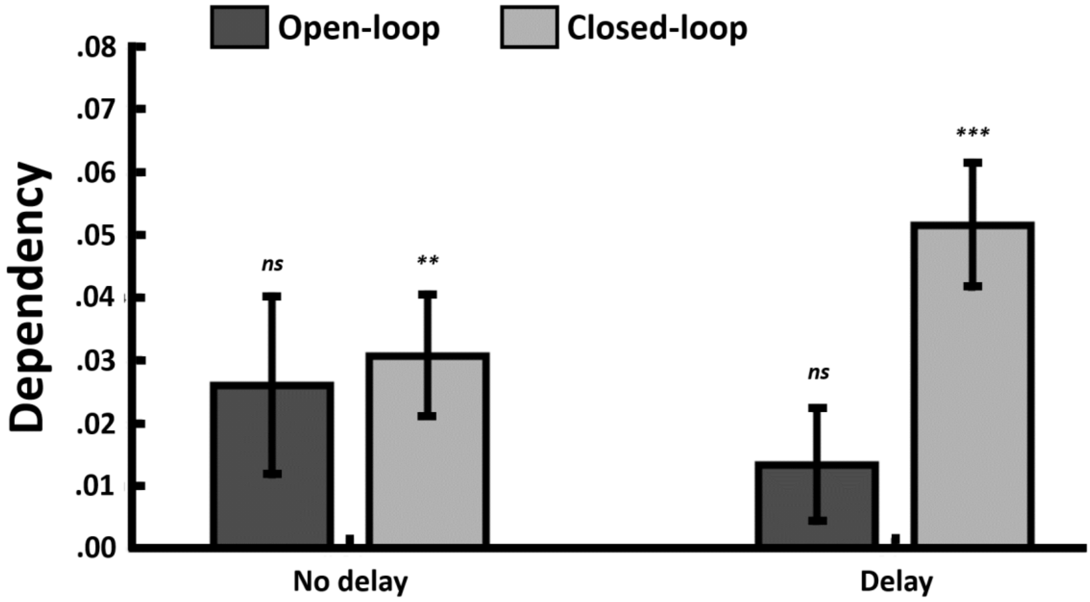

Figure 8 Dependency for open- and closed-loops at no-delay (i.e., encoded immediately prior to retrieval) and delay (i.e., encoded 24 hours prior to retrieval). Error bars represent ± 1 standard error. *ns* = not significant. *** $p < .001$, ** $p < .01$.

Retrieval dependency across element- and analysis-types. In order to assess whether differences in retrieval dependency between closed- and open-loops varied across element-types (i.e., people, locations, and objects) or analysis-type (i.e., common cue or target), I here report a 2x2x3 (Loop x Element-type x Analysis-type) ANOVA for dependency separately for the delay and no-delay conditions. Consistent with the analysis reported above, the ANOVA revealed no significant main effects or interactions at no-delay, Fs < 1.75, ps > .19. At delay, the ANOVA revealed a significant main effect of Loop, $F(1,29) = 9.38$, $p < .01$, $\eta_p^2 = .24$, consistent with the main effect of Loop seen in the dependency analysis (collapsed across element- and analysis-type) at delay reported above. No other significant main effects or interactions were seen, Fs < 3.07, ps > .09. Accordingly, I see no evidence for differences in retrieval dependency between closed-loops and open-loops as a function of the particular element-type (i.e., people, locations, or objects) or analysis-type (i.e., whether a particular element was a common cue or target)

4.4.2. fMRI

Reinstatement. To look for differences in reinstatement as an effect of Loop or Delay, I first identified regions where BOLD activity differed between the retrieval of different element-types (i.e., people, locations, and objects), collapsed across the closed- vs open-loop and delay vs no-delay conditions. For example, in order to identify regions associated with the retrieval of people, retrieval trials where a person was either a cue or target were compared against trials where a person was neither the cue nor target.

The largest difference in BOLD activity for people was seen in the precuneus, for locations in the left parahippocampal gyrus, and for objects in the left middle temporal gyrus (see Table 13). Thus, as in Horner et al. (2015), people, locations, and objects produced different levels of activation in specific neocortical regions, suggesting that different element-types are reinstated in distinct neocortical regions. BOLD responses from these regions were then extracted and assessed in terms of how they differed depending on whether the element associated with that region was either a cue, target, or non-target (e.g., a trial where an object was the cue and a location was the target would constitute a non-target trial for people). If closed-loops are reinstated holistically (regardless of delay), then the region associated with the non-target (i.e., the precuneus in the above example) should show greater activity for closed- relative to open-loops (as non-target reinstatement is not expected to occur for open-loops).

A 3x3x2 (Element-status x Region x Delay) ANOVA (with the dependent variable referring to differences in BOLD activity for closed- vs open-loops computed from the extracted BOLD responses) revealed a significant interaction between Element-status (i.e., cue vs target vs non-target) and delay (i.e., no-delay vs delay), $F(1,58) = 3.74$, $p = .03$, $\eta_p^2 = .11$.

Table 13 Clusters and peaks showing element-specific whole brain activity for people, locations, and objects (*p* < .05 FWE; cluster size > 30). L = Left hemisphere; R = Right hemisphere.

		MNI coordinates			
Region	**Voxels**	**X**	**Y**	**Z**	**Z score**
People					
Precuneus cortex	496	3	-54	27	> 8
L Lingual gyrus	1546	-12	-90	-3	> 8
Frontal medial cortex	273	-3	51	-9	> 8
L Middle temporal gyrus	184	-63	-6	-15	> 8
Superior frontal gyrus	214	-6	57	30	> 8
R Middle temporal gyrus	73	63	-3	-21	7.45
Locations					
L Parahippocampal gyrus	153	-27	-36	-15	> 8
L Precuneus cortex	107	-12	-54	12	> 8
R Precuneus cortex	99	15	-48	9	> 8
L Lateral occipital cortex	223	-39	-81	33	> 8
R Parahippocampal gyrus	66	27	-30	-18	7.50
Objects					
L Middle temporal gyrus	303	-54	-57	-3	7.33
L Frontal pole	114	-48	42	15	6.57
L Supramarginal gyrus	77	-60	-33	45	5.74
L Inferior frontal gyrus	35	-42	6	21	5.55

Further analysis revealed a significant effect of Element-status at delay, *F*(2,58) = 3.39, *p* = .04, η_p^2 = .11. I saw no evidence for an effect of Element-status at no-delay, *F*(2,58) = 1.37, *p* = .26, η_p^2 = .05. To interrogate this effect further, I compared BOLD activity in the closed- vs open-loop conditions separately for the delay and no-delay conditions. In both the delay and no-delay conditions, I assessed whether BOLD activity differed between closed and open-loops separately for cue, target, and non-target trials. At delay, closed-loops, relative to open-loops, showed greater activity during non-target trials, *t*(29) = 3.58, *p* < .01, *d* = .32, but not for cue, *t*(29) = 1.23, *p* = .23, *d* = .11, nor target trials, *t*(29) = 1.21, *p* = .24, *d* = .09. Activity for closed-loops did not differ significantly from open-loops for neither cue, target nor non-target trials at no-delay, *t*s < 1.51, *p*s > .14 (see Figure 9B).

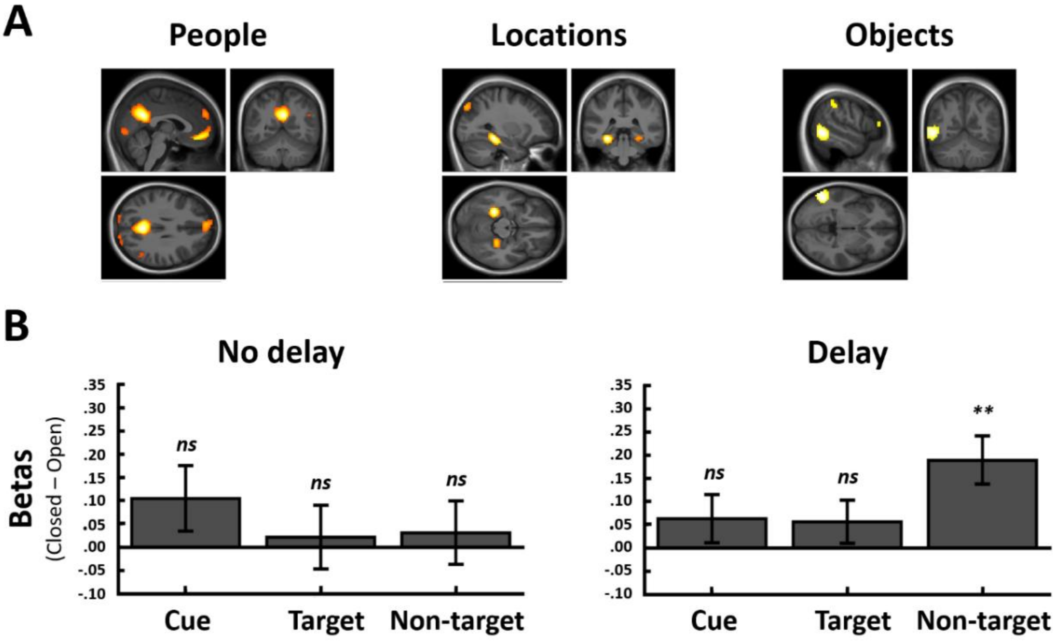

Figure 9 (A) Cortical regions showing activation differences for people, locations, and objects. **(B)** Mean difference (collapsed across precuneous, left parahippocampal cortex, and left middle temporal gyrus for people, locations, and objects, respectively) in BOLD activity for closed vs open-loops for cue, target and non-target trials at no-delay and delay. Error bars represent ± 1 standard error. *ns* = not significant. ** *p* < .01.

Horner et al., (2015) previously observed a non-target reinstatement effect for closed-loops retrieved immediately after encoding, and as such the failure to replicate this effect was not expected. However, the absence of a difference between closed- and open-loops at no-delay is perhaps consistent with the lack of a significant behavioural difference in dependency between closed- and open-loops at no-delay. Given that marginal evidence for dependency was seen for open-loops at no-delay, it is possible that this open-loop condition is obscuring a possible non-target reinstatement effect in the closed-loop condition. To account for this possibility, I compared BOLD activity for cue, target, and non-target trials for closed-loops at no-delay with BOLD activity for cue, target, and non-target trials for open-loops at delay. This is because no evidence for dependency was seen for open-loops retrieved after a delay, and as such they may act as a more appropriate baseline for assessing the presence of non-target reinstatement at no-delay (relative to open-loops at no-delay).

The comparison between closed-loops at no-delay vs open-loops at delay revealed that BOLD activity was significantly greater for closed-loops, relative to open-loops, for the non-target trials, $t(29) = 3.82$, $p < .001$, $d = .44$. No significant difference in BOLD activity was seen for target trials, $t(29) = 1.03$, $p = .31$, $d = .11$, but a difference was seen for cue trials, $t(29) = 2.33$, $p = .03$, $d = .22$ (see Figure 11A). This effect for cue trials did not survive correction for multiple comparisons. Thus, when using a baseline where no marginal evidence for dependency is seen for open-loops (i.e., open-loops at delay), I see evidence for non-target reinstatement for closed-loops in both the delay and no-delay condition. A comparison between this baseline estimate (i.e., closed-loops at no-delay vs open-loops at delay) and the difference between closed- and open-loops seen at delay (i.e., closed-loops at delay vs open-loops at delay), showed that BOLD activity associated with the reinstatement for non-target trials at delay did not differ significantly from that at no-delay, $t(29) = 1.11$, $p = .28$, $d = 21$.

The 3x3x2 (Element-status x Region x Delay) ANOVA also revealed a significant main effect of Region, $F(2,58) = 5.35$, $p < .01$, $\eta_p^2 = .16$. The factor Region refers to the three regions showing the largest BOLD response to each of the element-types (i.e., people, locations, and objects). These regions (i.e., precuneus, parahippocampal gyrus, and middle temporal gyrus) span superficial and deep cortical areas, and as such it is perhaps not surprising that the difference in BOLD activity between closed- and open-loops might differ across the three regions. In particular, regions within the medial temporal lobes, such as the parahippocampal gyrus, are sensitive to signal distortion and signal loss due to field inhomogeneities (Olman, Davachi, & Inati, 2009; Weiskopf, Hutton, Josephs, & Deichmann, 2006). Note that power calculations for the current sample size were estimated from Horner et al. (2015) where the difference in BOLD activity for closed and open-loops is collapsed across three regions of interest. It is possible that a larger sample size is needed to adequately assess differences in BOLD activity across the three regions used here (in particular regions sensitive to signal loss and distortion, such as the parahippocampal gyrus), as compared to when activity is collapsed across all three regions.

The ANOVA also revealed a significant three-way interaction between Element-status, Region, and Delay, $F(4,116) = 2.79$, $p = .03$, $\eta_p^2 = .09$. To interrogate this three-way interaction further, I performed a 2x2 (Element-status x Region) ANOVA separately for the delay and no-delay condition. At delay, the ANOVA revealed a significant main effect of Region, $F(2,58) = 3.93$, $p = .03$, $\eta_p^2 = .12$, with differences in BOLD activity for closed- and open-loops (collapsed across cue, target, and non-target trials) being significantly greater in regions associated with the reinstatement of people (i.e. precuneus), relative to locations (i.e., parahippocampal gyrus), $t(29) = 3.09$, $p < .01$, $d = .59$, and objects (i.e., medial temporal gyrus) relative to locations, $t(29) = 2.44$, $p = .02$, $d = .51$. Consistent with the main analysis, the 2x2 (Element-status x Region) ANOVA at delay revealed a significant effect of Element-status, $F(2,58) = 3.39$, $p = .04$, $\eta_p^2 = .11$. Critically, this effect of Element-status did not significantly interact with Region, $F(4,116) = 1.34$, $p = .34$, $\eta_p^2 = .04$.

At no-delay, the ANOVA also revealed a significant effect of Region $F(2,58) = 4.10$, $p = .02$, $\eta_p^2 = .12$, with a greater difference in BOLD activity for closed- and open-loops (collapsed across cue, target, and non-target trials) in the region associated with the reinstatement of objects compared to locations, *t*(29) = 2.36, *p* = .03, *d* = .48. A borderline effect was also seen between BOLD activity associated with the reinstatement of objects, relative to people, *t*(29) = 2.01, *p* = .05, *d* = .32. No other significant effects or interactions were seen, *F*s < 1.86, *p*s > .12.

Non-target reinstatement and correlated activity. The above analysis provided evidence that closed-loops retain their coherence over a 24-hour delay. Specifically, the analysis showed that non-target elements are reinstated during retrieval for closed-loops and that the level of reinstatement seen after a 24-hour delay does not differ from closed-loops retrieved immediately after encoding (when comparing closed-loops at no-delay vs open-loops at delay).

In previous work, Horner et al. (2015) showed that the strength of reinstatement during non-target trials correlated with hippocampal activity. To assess whether the role of the hippocampus may differ for closed-loops retrieved immediately after encoding and following a 24-hour delay, I looked for differences in activity, across participants, during retrievals for closed- and open-loop trials that correlated with differences between closed- and open-loops in cortical activity during non-target trials, separately for delay and no-delay condition.

Interestingly, while multiple regions that have previously been associated with recollection (see Kim, 2010; Rugg & Vilberg, 2013; Spaniol et al., 2009 for review) correlated with the difference in non-target reinstatement between closed- and open-loops at no-delay on a whole brain level, only the posterior left middle frontal gyrus (MFG) correlated with the non-target reinstatement effect at delay on a whole brain level (see Table 14).

Table 14 Clusters and peaks showing whole-brain regions correlating with differences in cortical activity for non-target trials between closed- and open-loops at delay and no-delay ($p < .05$ FWE, cluster size > 5). L = Left hemisphere; R = Right hemisphere.

		MNI coordinates			
Region	**Voxels**	**X**	**Y**	**Z**	**Z score**
Delay					
L Middle frontal gyrus	6	-39	-3	63	5.22
No-delay					
Paracingulate gyrus	159	9	45	9	5.96
R Cerebellum	27	27	-42	-42	5.50
R Angular gyrus	24	48	-57	15	5.34
L Lateral occipital cortex	79	-57	-69	-9	5.31
Cingulate gyrus	57	0	-39	42	5.27
R Lateral occipital cortex	16	45	-75	-15	5.14
L Paracingular gyrus	15	-15	42	-3	5.12
R Cerebellum	10	45	-57	-30	5.12
R Middle temporal gyrus	18	54	-9	-15	5.12
R Precuneus cortex	38	15	-51	12	5.05
Cerebellum	9	-9	-66	-36	5.02
L Middle temporal gyrus	13	-60	-36	-6	4.95
R Lingual gyrus	6	15	-48	0	4.92
R Superior temporal gyrus	8	54	-33	6	4.91
R Supramarginal gyrus	11	48	-42	12	4.87
Cingular gyrus	7	0	-6	45	4.82

Within the bilateral hippocampal mask, I saw that differences in BOLD activity for non-target trials correlated with hippocampal activity (consistent with Horner et al., 2015) at no-delay (see Table 15). However, no correlated activity in the hippocampus was seen at delay, suggesting that the role of the hippocampus in non-target reinstatement may differ for closed-loops retrieved immediately after encoding and those retrieved 24-hours after encoding.

Table 15 Clusters and peaks showing regions within the bilateral hippocampal mask correlating with differences in cortical activity for non-target trials between closed- and open-loops at no-delay ($p < .05$ FWE SVC, cluster defining threshold $p < .001$ uncorrected). L = Left hemisphere; R = Right hemisphere.

		MNI coordinates			
Region	**Voxels**	**X**	**Y**	**Z**	**Z score**
No-delay					
L Fusiform gyrus	4	-36	-18	-21	4.29
L Hippocampus	24	-12	-39	3	4.10
L Hippocampus	5	-33	-12	-12	3.92
L Fusiform gyrus	18	-39	-27	-15	3.83
R Parahippocampal gyrus	6	24	-39	-3	3.56
R Hippocampus	1	18	-36	0	3.53

To assess this possibility further, I identified regions in the hippocampus that showed greater BOLD activity to closed-, relative to open-loops (across all retrieval trials) at delay and no-delay. This analysis revealed a significant effect in the left hippocampus (peak: -21, -36, 6) for the delay condition. Note that the same left hippocampal effect (peak: -21, -36, 6) was seen for the no-delay condition, but only at a lenient uncorrected threshold ($p < .05$). This is perhaps not unexpected given the absence of a difference in BOLD activity between closed- and open-loops retrieved immediately after encoding, however the purpose of this analysis was only to identify voxels in the hippocampus, in a relatively non-biased manner (in contrast to the hippocampal regions identified in the above analysis (see Table 15)), to then correlate with non-target reinstatement in the neocortex.

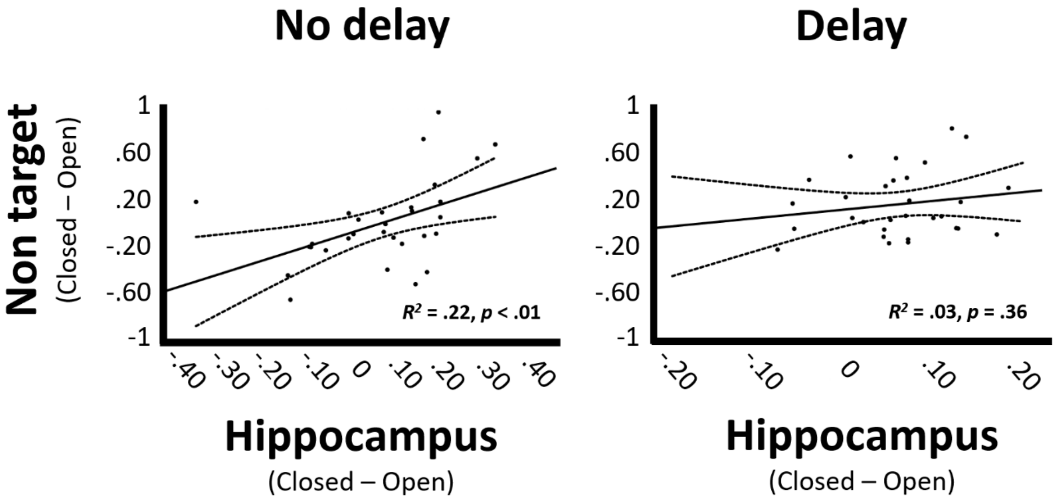

Figure 10 Correlation between the difference in hippocampal activity during the retrieval of closed- vs open-loops (hippocampal ROI defined as a 9-mm radius sphere centred on the peak coordinate from the closed- vs open-loop contrast; peak: -21, -36, 6) and the mean difference (collapsed across precuneous, left parahippocampal cortex, and left middle temporal gyrus for people, locations, and objects, respectively) in BOLD activity during non-target retrieval for closed vs open-loops.

Consistent with the analysis reported above, I saw that activity from the closed- vs open-loop contrast in the hippocampus (i.e., mean activity per participants) correlated with participants' non-target reinstatement at no-delay, r = .47, p < .01. No significant correlation between hippocampal activity and non-target reinstatement was seen at delay, r = .17, p = .36 (see Figure 10). Importantly, a significant correlation was also seen between hippocampal activity and non-target reinstatement when contrasting closed-loops at no-delay with open-loops at delay, r = .63, p < .001 (see Figure 11B) suggesting that the correlation seen at no-delay is not an artefact of the particular open-loop baseline used (i.e., open-loops at delay vs open-loops at no-delay).

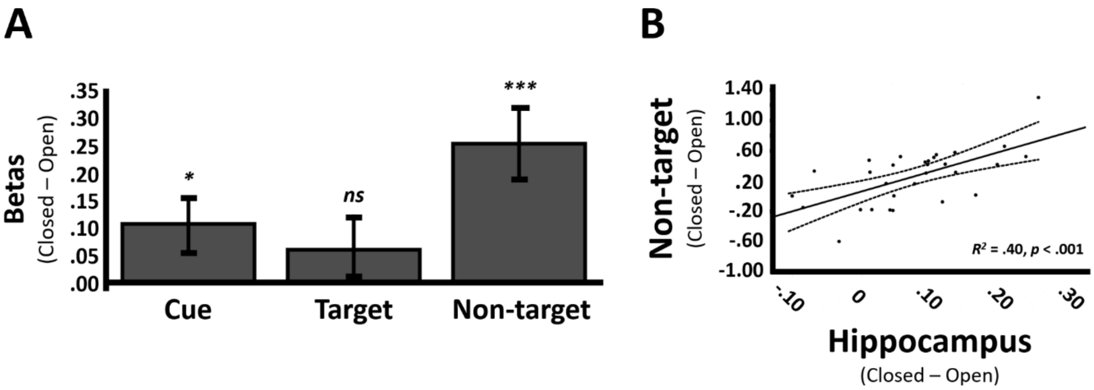

Figure 11 (A) Mean difference (collapsed across precuneous, left parahippocampal gyrus, and left middle temporal gyrus for people, locations, and objects, respectively) in BOLD activity for closed-loops at no delay vs open-loops at delay for cue, target and non-target trials. Error bars represent ± 1 standard error. *ns* = not significant. * $p < .05$ (uncorrected), *** $p < .001$. **(B)** Correlation between the difference in hippocampal activity during the retrieval of closed-loops at no-delay vs open-loops at delay (hippocampal ROI defined as a 9-mm radius sphere centred on the peak coordinate from the closed- vs open-loop contrast; peak: -21, -36, 6) and the mean difference (collapsed across precuneous, left parahippocampal cortex, and left middle temporal gyrus for people, locations, and objects, respectively) in BOLD activity during non-target retrieval for closed-loops at no-delay vs open-loops at delay.

A comparison between the correlation observed at delay and no-delay revealed that the correlation between the closed- vs open-loop contrast in the hippocampus and participants' non-target reinstatement effect in the neocortex did not differ significantly between the delay and no-delay conditions, $Z = 1.24$, $p = .21$. Note though that the comparison of the correlation between hippocampal activity and non-target reinstatement at delay and the correlation seen when contrasting closed-loops at no-delay with open-loops at delay revealed a significant difference between the two correlations, $Z = 2.09$, $p = .04$, but this correlation did not survive correction for multiple comparisons.

4.5. General discussion

In Chapter 2, I provided behavioural evidence that closed-loops are associated with retrieval dependency after a period of forgetting. Here, I provide fMRI evidence in support of this and show that closed-loops continue to be retrieved holistically following a 24-hour interval between encoding and retrieval. Retrieval dependency was seen for closed-loops retrieved immediately after encoding and, critically, following a 24-hour delay. Consistent with this behavioural evidence, neocortical reinstatement of elements incidental to retrieval was observed for closed-loops retrieved 24 hours after encoding. For example, when cued with a location (e.g., *kitchen*) and asked to retrieve an object (e.g., *hammer*), activity was seen in cortical regions 'representing' the associated person (e.g., *Barack Obama*). The level of non-target reinstatement seen at delay was equivalent to that seen for closed-loops retrieved immediately after encoding (when comparing closed-loops at no-delay vs open-loops at delay), providing support for the behavioural evidence presented in Chapter 2 suggesting that closed-loops retain their coherence following a period of forgetting.

Reinstatement of non-target elements is thought to be a consequence of pattern completion in the hippocampus (Horner & Doeller, 2017). Consistent with this, Horner et al. (2015) previously showed that differences in neocortical reinstatement of non-target elements between closed- and open-loops correlated with differences in activity in the hippocampus. Furthermore, recent evidence has shown that the correlation between non-target reinstatement and the hippocampus involves the CA3 subregion of the hippocampus (Grande et al., 2019), consistent with the proposed role of CA3 in pattern completion (Neunuebel & Knierim, 2014). Here, I observed that differences between closed- and open-loops in the reinstatement for non-target elements retrieved immediately after encoding correlated with differences in activity between closed- and open-loops in the hippocampus; providing further support for the role of the hippocampus in the holistic retrieval of event-based representations (Cohen & Eichenbaum, 1993; Eichenbaum, Yonelinas, & Ranganath, 2007; Mayes et al., 2007). The involvement of the hippocampus in neocortical reinstatement for closed-loops – in the absence of a significant difference when comparing non-target trials for closed- and open-loops at no-delay (rather than closed-loops at no-delay and open-loops at delay) – suggests that the overlapping associative structure formed during the encoding of closed-loops differs from the associations formed during the encoding of open-loops, even when marginal behavioural evidence for dependency is seen for open-loops.

It is possible that mnemonic generalisation (Shohamy & Wagner, 2008), or reactivation during encoding (Zeithamova et al., 2012) of two overlapping associations, encoded as open-loops, can lead to increases in dependency. For example, when a participant encodes the associations *kitchen-Barack Obama*, and *Barack Obama-hammer*, the reinstatement of *Barack Obama* during the encoding of *Barack Obama* and *hammer* may strengthened the association between *kitchen and hammer*, leading to dependency in the absence of explicitly encoding all the possible associations. It is also possible that changes in encoding strategies, when the associative structure of the overlapping pairs becomes apparent to the participants, can lead to dependency for open-loops. It has been argued that the use of within-subjects designs can lead to transfers of encoding strategies that are optimal for one condition, but not another, when conditions are interleaved randomly (as in the current experiment) (Poulton, 1982). Participants have perhaps inferred the associative nature of the overlapping pairs following the first encoding session (i.e., delay) and used a different encoding strategy for the second encoding session (i.e., no-delay) leading to changes in dependency for both open- and closed-loops.

However, consistent with Horner et al. (2015), the correlation between differences in non-target reinstatement and hippocampal activity between closed- and open-loops observed here, even when BOLD activity at retrieval does not differ significantly between closed- and open-loops (when comparing non-target trials for closed- and open-loops at no-delay), suggests that the explicit encoding of all possible associations in a closed-loop associative structure may be a necessary boundary condition for the occurrence of pattern completion (Horner et al., 2015). For instance, Horner et al., (2015) showed that activity in the hippocampus, specifically during the encoding of the third, and final pair (e.g., *kitchen-hammer*) predicted performance for the other two associations (e.g., *kitchen-Barack Obama* and *Barack Obama-hammer*) in the closed-loop condition, but not for open-loops consisting of an associative chain of three pairwise associations (e.g., *kitchen-hammer*, *kitchen-Barack Obama*, *Barack Obama-dog*). This is consistent with evidence suggesting that activity in the hippocampus immediately following the offset of a naturalistic ‘event’ may reflect the integration of event elements into a coherent event representation (Ben-Yakov & Dudai, 2011).

No significant correlation between hippocampal activity and the neocortical reinstatement of non-target elements was observed for closed-loops retrieved following a 24-hour delay. Note though that the correlation between hippocampal activity and neocortical reinstatement of elements incidental to retrieval did not differ significantly for closed-loops retrieved immediately after encoding and following a 24-hour delay. Interestingly, the absence of a significant correlation was seen despite evidence for non-target reinstatement occurring at delay. This observation is consistent with previous findings showing that the role of the hippocampus in retrieval of episodic memories decreases following a period of consolidation (Niki & Luo, 2002; Piefke et al., 2003; Watanabe et al., 2012). The sparse activity patterns and high neural density of the hippocampus are thought to provide the ideal mechanism for the initial acquisition and retrieval of an event-based representation (McClelland et al., 1995; McNaughton & Morris, 1987; Norman & O’Reilly, 2003; Treves & Rolls, 1994). However, over a period of consolidation, gradual adjustments to synaptic connections between cortical neurons representing elements from the same event are thought to allow for the formation of a neocortical representation that can be retrieved independently of the hippocampus (McClelland et al., 1995). Consistent with this, previous work has shown that connectivity between neocortical regions representing two associated elements increases over a 24-hour interval that includes a night of sleep (Takashima et al., 2009)

McClelland et al. (1995) (but see Kumaran, Hassabis, & McClelland, 2016) have argued that this consolidation process is relatively protracted with differences in hippocampal activity between recent and remote memories only emerging after days, weeks, or months. However, evidence from a rodent study has shown that the consolidation of an initially hippocampal-dependent memory representation can occur rapidly when memory formation is supported by pre-existing knowledge (Tse et al., 2007). Tse et al. (2007) showed that rats were able to learn two flavour-place associations following a single exposure when the novel associations were supported by flavour-place associations learnt previously in the same spatial context. Importantly, these associations persisted in memory even when the hippocampus was surgically removed two days after learning, suggesting that recently formed memory representations can rapidly become independent of the hippocampus if they are supported by pre-existing knowledge. This is consistent with evidence showing that patients with amnesia can more readily learn the associations between household items and their prices if the prices in question are congruent with the patients pre-existing knowledge of an items ‘market value’ (Kan, Alexander, & Verfaellie, 2009).

The paradigm used here requires participants to encode triplets consisting of famous people, locations, and common objects. Thus, it is possible that participants can draw on their pre-existing knowledge of the elements during encoding. As such, the retrieval of closed- and open-loops may not only require the reinstatement of the specific elements associated with a particular ‘event’, but also participants’ conceptual knowledge recruited during the encoding of the pairwise associations (Renoult, Irish, Moscovitch, & Rugg, 2019). It is possible that participants’ pre-existing knowledge about the elements and/or the congruency of the pairwise associations may accelerate the rate of consolidation (e.g., Tse et al., 2011). Although the elements used in the current experiment were well-known examples of famous people, locations, and common objects, the people-location-object triplets were randomly created for each participant. Further work is needed to assess whether differences in the conceptual congruency of individual triplets may lead to differences at immediate and delayed retrieval on both a behavioural and neural level.

The findings presented here provided evidence to support standard consolidation theories (e.g., McClelland et al., 1995; Squire & Alvarez, 1995). However, it is less clear what the absence of a significant correlation between hippocampal activity and neocortical reinstatement might imply for the retrieval of the underlying memory representation. At encoding, participants are explicitly instructed to try and imagine the associations interacting in a vivid and meaningful way in order to encourage deeper and more elaborative encoding (Paivio, 1969). It is possible that the absence of a significant correlation between hippocampal activity and the neocortical reinstatement of non-target elements at delay might reflect a 'loss' of the elaborative context established at encoding (e.g., Addis et al., 2004; Geib, Stanley, Wing, Laurienti, & Cabeza, 2017; Gilboa et al., 2004). This is in light of the proposal that the involvement of the hippocampus is necessary for reinstating event-based representations in their original and rich contextual form (Winocur & Moscovitch, 2011). As such, the elements that make up a closed-loop can still be reinstated in their entirety following a delay - without the involvement of the hippocampus - but the contextual details of the 'event' might be forgotten. This proposal would be consistent with evidence showing that peripheral details of an event are forgotten more rapidly than central, gist-like details (Sekeres et al., 2016), and evidence showing that memory for contextual details surrounding an event tends to be impaired following hippocampal damage (St-Laurent, Moscovitch, Levine, & McAndrews, 2009; Steinvorth, Levine, & Corkin, 2005). This proposal would, nevertheless, be premature given the evidence presented here. Further work is needed to reveal whether the reinstatement of elements belonging to closed-loops retrieved immediately after encoding and following a 24-hour delay differ in some qualitative manner. It is also important to note that the correlation between hippocampal activity and neocortical activity associated with the reinstatement of non-target elements did not differ significantly between closed-loops retrieved immediately after encoding and following a 24-hour delay. Thus, the extent to which strong claims can be made about the absence of a significant correlation between hippocampal activity and neocortical reinstatement following a delay is limited.

On a whole brain level, regions within the default mode network (Buckner, Andrews-Hanna, & Schacter, 2008) correlated with the reinstatement of non-target elements retrieved immediately after encoding. In light of the proposed role of the default mode network in ‘scene construction’ (see Hassabis & Maguire, 2007 for review) and self-generated thought (Mason et al., 2007), it is possible that the correlation between the regions observed here and non-target reinstatement may reflect an internalised process at retrieval that reconstructs the ‘event’ formed at encoding, independently of the explicit demands of the cued-recognition task used here. This would be consistent with the proposal that the default mode network may support recollection via a process that is ‘constructive’ in nature (Addis, 2018). Note though that some regions that showed correlated activity with differences in non-target reinstatement for closed- and open-loops overlap with regions showing BOLD activation differences for the different element-types (i.e., people, locations, and objects), and as such it is difficult to disentangle the effects of the retrieval of the elements themselves and any internalised mental activity at retrieval.

Only the left MFG correlated with differences in non-target reinstatement for closed- and open-loops retrieved 24-hours after encoding. The left MFG has previously been associated with retrieval effort (Buckner, Koutstaal, Schacter, Wagner, & Rosen, 1998) and retrieval monitoring processes (Achim & Lepage, 2005; Dobbins, Foley, Schacter, & Wagner, 2002). However, the coordinates reported here refer to a more posterior portion of the MFG that overlaps with the precentral gyrus. It is therefore not clear whether the activation observed here could be considered functionally equivalent to that reported in previous work. Further work is needed to assess the exact contributions of regions identified on the whole brain level to the reinstatement of elements incidental to retrieval, both immediately after encoding and following a delay.

4.6. Conclusion

I provided evidence for holistic reinstatement for closed-loops retrieved following a period of forgetting and consolidation. This is consistent with behavioural findings reported in Chapter 2, suggesting that event-based representations retain their coherence following a period of forgetting. However, despite the proposed role of the hippocampus in episodic memory retrieval (Aggleton & Brown, 1999; Eichenbaum et al., 2007; Mayes et al., 2007), I observed no (significant) evidence that the 'incidental' reinstatement of closed-loops elements was supported by the hippocampus following a period of forgetting and consolidation. This finding provides some support for the proposal that the role of the hippocampus in recollection may decrease following a period of consolidation (McClelland et al., 1995; Squire & Alvarez, 1995).

4.7. Acknowledgments

I wish to thank the staff at the York Neuroimaging Centre (YNiC) for their assistance in running this experiment. I also wish to thank Dr Sam Berens and Dr André Gouws for their help with piloting the fMRI acquisition.

5. General discussion

5.1. Summary of findings

In this thesis, I aimed to assess how event-based memories changed over a period of consolidation and forgetting. I used a paradigm developed in Horner and Burgess (2014) that allows for forming ‘events’ that are matched in the number of elements but differ in their coherence at retrieval, depending on the associative structure formed at encoding (i.e., closed- vs open-loops). The experiments presented in this thesis showed that the associative structure formed at encoding has a lasting impact on the coherence at retrieval.

5.1.1. Chapter 2

In Chapter 2, I used retrieval accuracy and dependency to assess how event-based representations change over a period of forgetting. I reasoned that if forgetting (as reflected in decreases in retrieval accuracy) is associated with decreases in dependency, then this would imply that closed-loops fragment as a function of forgetting. However, if forgetting is occurring, but dependency remains consistent over time, then this would suggest that closed-loops tend to be forgotten in an all-or-none manner, such that closed-loops are more likely to be either retained or forgotten in their entirety.

In Experiment 1, I assessed retrieval accuracy and retrieval dependency immediately following encoding and a 12-hour delay. Despite evidence for forgetting, I saw no evidence for changes in dependency over a 12-hour interval. In Experiment 2, I increased the length of the interval between immediate and delayed retrieval to one week. Although I saw numerical increases in forgetting in Experiment 2, relative to Experiment 1, I observed no evidence for significant decreases in dependency as a function of forgetting.

In Experiments 1-2, participants encoded either closed- *or* open-loops. Although it has previously been shown that closed-loops are associated with greater retrieval accuracy than open-loops (Horner et al., 2015), I saw little evidence for differences in retention for closed- and open-loops. In Experiment 3 and 4, I assessed whether retrieval accuracy differed between closed- and open-loops when participants learnt both closed- *and* open-loops (i.e., a within-subject design, compared to the between-subject design used in Experiments 1-2). Experiments 3-4 also offered an opportunity to further assess whether retrieval dependency decreased following a period of forgetting. In Experiment 3, I provided evidence that closed-loops were remembered better than open-loops following a one-week delay. In Experiment 4, I saw that closed-loops were also associated with greater retrieval accuracy immediately after encoding but forgetting rates for closed-loops were also significantly lower than for open-loops over a week-long delay. Further, consistent with Experiment 1-2, I saw no evidence for changes in dependency for closed-loops (or open-loops) as a function of forgetting.

In Chapter 2, I provided consistent evidence that dependency for closed-loops does not change as a function of forgetting, even in instances where forgetting was high. This suggests that the individual elements or associations that make up a closed-loop are forgotten at relatively similar rates, at least so far as the rate of forgetting is sufficiently uniform for closed-loops to retain their coherence after a period of forgetting. Further, I observed that the associative structure formed at encoding led to differences in retention, such that closed-loops were better retained than open-loops; but only when participants learnt both closed- *and* open-loops in a within-subject design.

5.1.2. Chapter 3

In Chapter 2, I also assessed whether sleep supports mnemonic integration for overlapping associations, under the assumption that if two overlapping associations (encoded in the open-loop condition) are integrated during sleep, then dependency should be seen. In Chapter 2, I observed that sleep improved inference for non-encoded pairs but did not increase dependency for the directly encoded associations. This suggests that the role of sleep in supporting inference is unlikely to be driven by the integration of two overlapping associations into a generalised representation, as proposed by Lau et al. (2010). Sleep is perhaps more likely to support inference via a dynamic, ‘on-the-fly’ process at retrieval (Kumaran & McClelland, 2012); the probability of which is increased due to decreases in forgetting for directly encoded associations following sleep, relative to wakefulness.

In Chapter 3, I used a technique called targeted memory reactivation (TMR) (Rudoy et al., 2009) to induced memory reactivation of open-loops during sleep. This was done in order to assess whether systems consolidation as a function of memory reactivation during sleep, rather than sleep more generally (as in Chapter 2) supports mnemonic integration for open-loops. If memory reactivation during sleep supports a process where two overlapping associations are integrated into a generalised representation (Lau et al., 2010), then retrieval dependency should be seen following sleep, but only for open-loops associated with sound cues that participants were re-exposed to during sleep. Despite previous evidence for increased inference following sleep (Lau et al., 2010) and proposals suggesting that sleep strengthens associations between overlapping memories (Diekelmann & Born, 2010), I saw no evidence for increases in inference as a function of TMR, nor was TMR associated with increases in dependency. Instead, I provided evidence that TMR during sleep led to decreases in retention for the second (or weakly) encoded association of an open-loop. I also observed evidence for improved retention for the first (or strongly) encoded association following TMR, but this effect did not survive correction for multiple comparisons.

Consistent with previous work (Oyarzun et al., 2017; Antony et al., 2018), Chapter 3 provided evidence for forgetting following TMR, possibly as an effect of proactive interference during encoding. It is possible that variations in encoding strength led to differences in the level of activation for the first and second encoded associations during TMR. Further work is needed to assess this possibility, but given this assumption, the findings presented in Chapter 3 are consistent with retrieval-induced forgetting effects seen during wakefulness (Anderson et al., 1994; Wimber et al., 2015), where repeated retrievals of a 'target' memory can lead to decreases in retention for a 'competing' memory representation (see Anderson, 2003 for review).

5.1.3. Chapter 4

In Chapter 4, I used fMRI to examine whether closed-loops are retrieved holistically following a period of forgetting and consolidation. Previous work has shown that closed-loops are associated with the reinstatement of all event elements, including those incidental to the specific retrieval task, when retrieval follows immediately after encoding (Horner et al., 2015). Following the findings in Chapter 2 suggesting that closed-loops are associated with dependency following a period of forgetting, I assessed whether closed-loops retrieved following a 24-hour delay continue to be associated with the neocortical reinstatement of elements incidental to retrieval.

In previous work, Horner et al. (2015) showed that the strength of this 'incidental' reinstatement effect was correlated with activity in the hippocampus, for closed-loops retrieved immediately after encoding. However, theories of systems consolidation (e.g., McClelland et al., 1995; Nadel & Moscovitch, 1997; Squire & Alvarez, 1995; Winocur & Moscovitch, 2011) make different predictions about the role of the hippocampus in retrieval following a period of consolidation. Therefore, I also examined whether the reinstatement of elements incidental to retrieval correlated with activity in the hippocampus for closed-loops retrieved immediately after encoding (as in Horner et al., 2015) and following a 24-hour delay.

I provided evidence that the retrieval of closed-loops following a period of forgetting and consolidation was associated with neocortical activity corresponding to event elements incidental to the specific retrieval task. However, while activity in the hippocampus for closed-loops retrieved immediately after encoding correlated with the reinstatement of closed-loop elements incidental to retrieval, I did not find evidence for a significant correlation between this ‘incidental’ reinstatement effect and hippocampal activity for closed-loops retrieved following a 24-hour delay.

Consistent with behavioural evidence from Chapter 2, the findings presented in Chapter 4 suggested that closed-loops continue to be retrieved in a holistic (and coherent) manner after a period of forgetting and consolidation. The findings also provided some support for the proposal that the role of the hippocampus in retrieval may decrease over a period of consolidation (McClelland et al., 1995; Squire & Alvarez, 1995), but note that the correlations immediately after encoding and following a delay did not differ significantly from each other, and as such the extent to which strong claims can be made about the absence of a significant correlation following a delay is limited.

5.2. All-or-none forgetting

In Chapter 2 I demonstrated that closed-loops are consistently associated with dependency, irrespective of whether retrieval follows immediately after encoding or after a 12-hour or one-week delay. This was taken as evidence that closed-loops tend to be forgotten in an all-or-none manner, such that closed-loops are more likely to be either retained or forgotten in their entirety.

However, any measure of forgetting will inevitably be derived from differences in the proportion of ‘events’ retained between two retrieval time points. Thus, the proposal that forgetting is all-or-none necessitates making an assumption about the coherence of closed-loops that are forgotten. Closed-loops that fragment as a function of forgetting (where some pairwise associations are retained, and others are forgotten within the same closed-loop) would increase the proportion of ‘events’ that fall along the off-diagonal cells in the contingency tables used to calculate the proportion of joint retrieval for the data and independent model. This would consequently produce levels of dependency that are less than that seen for closed-loops retrieved immediately after encoding. To quantify this assumption, I compared the level of dependency seen for closed-loops retrieved after a period of forgetting with a model of independent forgetting and showed that the observed dependency is greater than that predicted under an assumption of independent forgetting. While this does not exclude the possibility that some closed-loops might fragment over a period of forgetting, it suggests that closed-loops are more likely to be forgotten in an all-or-none manner.

Note though that the method for calculating dependency necessitates that dependency is estimated *across*, rather than *within* events. I refer to the 'raw' measure of dependency as the proportion of joint retrieval but note that this measure also includes the proportion of joint non-retrieval. As the raw measure of dependency is calculated separately for each of the three common cue (e.g., cue with *Obama*, retrieve *kitchen* or *hammer*) and common target (e.g., cue with *kitchen* or *hammer*, retrieve *Obama*) combinations, there could, for example, be instances where a closed-loop is associated with joint retrievals for certain common cue or target combinations, but joint non-retrievals for other combinations. Because dependency is averaged across all closed-loops separately for each contingency table, a closed-loop that is associated with a proportion of joint retrievals for some cue-target combinations and joint non-retrieval for other cue-target combinations will contribute equally to dependency as a closed-loop where cue-target combinations are all retrieved correctly or incorrectly. As such, a closed-loop can be retrieved in a relatively asymmetrical way, but still contribute to dependency in a symmetrical manner.

Although this pattern of retrieval is theoretically possible, the retrieval of individual associations has long been thought to be symmetrical, such that the likelihood of retrieving, for example, *kitchen* when cued by *Barack Obama* is related to the likelihood of retrieving *Barack Obama* when cued by *kitchen* (Asch & Ebenholtz, 1962). As such, it is unlikely that a particular closed-loop association will contribute to a leading diagonal cell (e.g., both correct) in one instance, but to the other leading diagonal cell in another instance (e.g., both incorrect). Further, the dependency measure used here is derived from previous work showing that retrieval dependency is consistently able to differentiate between the retrieval for closed- and open-loops (Horner et al., 2015; Horner & Burgess, 2013, 2014). Thus, it is difficult to envision why this potential confound should systematically produce evidence for dependency for closed-, but not open-loops.

5.2.1. All-or-none vs independent forgetting

A similar approach to the one presented in Chapter 2 was taken by Brady et al. (2013) who assessed dependency for different features (i.e., exemplar and state) of an object. However, in contrast to evidence for retained dependency for closed-loops, Brady et al. (2013) found that dependency decreased over time. This suggests that features associated with an object are forgotten in an independent manner (but see Balaban et al., 2019). However, evidence from the source-memory literature has suggested that there is a degree of asymmetry in relation to how source details (e.g., colour and location) are bound to items, with information about colour and location being bound to the item, but not each other (Starns & Hicks, 2008). This may underlie the decreases in dependency seen in Brady et al. (2013). As features of an item are represented in a relatively independent manner, they are perhaps also likely be to forgotten in a similarly independent way (but see Balaban et al., 2019). In contrast, elements of a closed-loop are proposed to be bound into a coherent event representation (Horner et al., 2015; Horner & Burgess, 2014) and as such the forgetting of elements and/or associations from a closed-loop are perhaps more likely to be related.

It is possible that differences in the pattern of forgetting observed here and in Brady et al. (2013) reflect some underlying differences between item- and event-based representations. Whereas item-based representations are thought to be supported by the perirhinal cortex

(Aggleton & Brown, 1999; Diana et al., 2007), the hippocampus has consistently been implicated in the binding processes required to form a multi-element event representation (Cohen et al., 1999; Davachi, 2006; Mayes et al., 2007; Ranganath, 2010) Critically, while neural representations in the perirhinal cortex are thought to be distributed and overlapping, the high neural density and sparse representations in the hippocampus (Barnes et al., 1990; Viskontas et al., 2006) are likely to reduce any representational overlap (Marr, 1971; McClelland et al., 1995; Norman & O'Reilly, 2003). This has led to the proposal that retrieval processes thought to be supported by the hippocampus (i.e., recollection) are less susceptible to forgetting via interference, relative to those thought to be supported by the perirhinal cortex (i.e., familiarity) (Sadeh et al., 2014). Instead, forgetting for hippocampal representations has been proposed to be a result of decay (Frankland et al., 2013; Hardt et al., 2013). The findings presented in Chapter 2 appear most consistent with this proposal and provided evidence to suggest that the rate of decay within an event is relatively uniform.

An interesting possibility for future work would be to assess whether the pattern of forgetting observed here for closed-loops may differ between the elements and features associated with the individual elements. Given the findings presented in Chapter 2, and those reported by Brady et al. (2013), it is possible that while the association between the elements themselves may be forgotten in an all-or-none manner, features associated with the elements may be forgotten independently of each other.

5.2.2. Continuous forgetting

I have made a distinction between all-or-none and independent forgetting. However, this does not exclude the possibility that closed-loops undergo a continuous form of forgetting. If the forgetting of event-based representations occurs via decay (Frankland et al., 2013; Hardt et al., 2013; Sadeh et al., 2014), then this would suggest that the forgetting of closed-loops is a gradual, continuous process that takes place over time. Yet, it is possible that the behavioural consequence of decay is relatively binary, in so far that the rate of decay for closed-loops will either be more or less than that required to ensure its 'accessibility' after a period of forgetting.

I have suggested that the pattern of forgetting for closed-loops reflects a process of decay that is relatively uniform for separate elements of an event. Although there may be variations in the amount of decay within an event, differences in decay should be greater *across*, relative to *within* events. Closed-loops perhaps undergo a continuous form of forgetting, but as long as the rate of decay within an event is sufficiently uniform for event elements to remain either above or fall below a given retrieval threshold (e.g., Norman & O'Reilly, 2003; Yonelinas, 1994), the behavioural consequence of decay should reflect a pattern of forgetting that is all-or-none.

However, recent work has suggested that recollection can be characterised by two, possibly separable memory components; one that characterises the 'accessibility' of the memory, and another that characterises the 'precision' of the memory (Harlow & Donaldson, 2013; Harlow & Yonelinas, 2016). Support for the proposal that the accessibility and precision of recollection are dissociable (Harlow & Yonelinas, 2016) has come from work showing that the success and precision of recollection are differentially affected by retrieval practice, with retrieval practice enhancing the accessibility, but not the precision, of later retrieval (Sutterer & Awh, 2016). However, a more recent study provided evidence to show that memory precision also benefits from retrieval practice (Schuetze, Eglington, & Kang, 2019). The accessibility and precision of recollection have also been shown to be dissociable on a neural level, with activity in the hippocampus associated with instances where recollection succeeds, and activity in the angular gyrus scaling with retrieval precision (Richter, Cooper, Bays, & Simons, 2016). Note though that a recent study showed that the precision of recollection was impaired in patients with hippocampal damage (Nilakantan, Bridge, Van Haerents, & Voss, 2018), suggesting that the hippocampus also contributes to the precision of memory retrieval.

Interestingly, Cooper and Ritchey (2019) recently demonstrated that the likelihood of successfully retrieving, for instance, the association between a spatial context and a colour of an object showed evidence for dependency. However, the precision with which the context and colour were retrieved showed no evidence for retrieval dependency. This suggests that although the elements that make up an 'event' may be retrieved in an all-or-none manner, the precision with which each of the individual elements are retrieved will differ in some independent way, even when retrieval follows shortly after encoding. As such, it is possible that closed-loops that retain their accessibility may differ in some continuous and qualitative manner before and after a period of forgetting.

5.2.3. Summary

In Chapter 2, I provided evidence to show that closed-loops are likely to be forgotten in an all-or-none manner. However, this does not exclude the possibility that the elements or associations themselves undergo some continuous form of forgetting over time. Theoretical accounts have proposed that different types of mnemonic information are forgotten at different rates (Reyna & Brainerd, 1995; Winocur & Moscovitch, 2011), and I have argued that the pattern of forgetting observed in Chapter 2 possibly reflects a rate of decay that is sufficiently uniform for closed-loops to either retain their accessibility or not. However, for closed-loops that 'survive' forgetting, the elements or associations that make up a closed-loop may differ in a continuous and independent manner before and after a period of forgetting.

5.3. Forgetting during sleep

In Chapter 2, I argued that the forgetting of closed-loops is consistent with a process of memory decay, where the rate of forgetting is relatively uniform within an event. However, in Chapter 3, I demonstrated that the forgetting of open-loops may, at least in some instances, be much less symmetrical. Specifically, I provided evidence that memory reactivation during sleep, as a function of TMR, led to forgetting for the second encoded association of an open-loop, possibly as a function of proactive interference at encoding. This pattern of forgetting is similar to retrieval-induced forgetting effects seen during wakefulness (see Anderson, 2003 for review), under the assumption that TMR is more likely to be associated with higher levels of reactivation for the first (given that memory performance was greater for this association prior to sleep), relative to the second encoded association of an open-loop.

5.3.1. Retrieval-induced forgetting during sleep

Retrieval-induced forgetting during wakefulness has been attributed to inhibitory processes established at retrieval that ‘suppress’ a competing memory, which may otherwise interfere with the retrieval of a ‘target’ memory (Kuhl et al., 2007; Wimber et al., 2015). However, I observed that a similar effect can emerge following repeated memory reactivation during sleep, presumably in the absence of any conscious executive control. This is consistent with findings previously reported by Oyarzún et al. (2017). However, given the emphasis on the role of inhibitory mechanisms in retrieval-induced forgetting (Anderson, 2003), it is less clear how such an effect may emerge during sleep; where inhibitory control during memory reactivation is thought to be present to a significantly lesser extent, relative to memory retrieval during wake (Hasselmo, 1999).

The findings presented in Chapter 3 appear most consistent with the proposal that differences in co-activation for two overlapping (or competing) memories can lead to retrieval-induced forgetting (Norman et al., 2007). Note that here I am assuming that differences in encoding strength between two overlapping open-loop associations increased the likelihood of the first encoded association being reinstated more strongly, relative to the second encoded association, as a function of TMR. Further work employing multivariate pattern analysis (MVPA) to EEG could address this possibility (e.g., Bramão & Johansson, 2018).

Computational accounts assume that forgetting occurs when the level of activation differs between two overlapping associations, such that one is reinstated strongly and the other only moderately (Newman & Norman, 2010; Norman et al., 2006, 2007). This is consistent with evidence showing that moderate, but not high or low, levels of activation of a 'to-be-suppressed' item in a think/no-think paradigm can lead to forgetting (Detre, Natarajan, Gershman, & Norman, 2013; see also Newman & Norman, 2010). Retrieval-induced forgetting is not thought to occur when two competing memories are both reinstated at relatively high levels, nor when the level of activation for a 'competitor' is particularly low. In case of the former, it is possible that the concurrent and strong reinstatement of overlapping memories may underlie the mnemonic integration processes associated with systems consolidation during sleep (Lewis & Durrant, 2011; Stickgold & Walker, 2013; Diekelmann & Born, 2010). Note also that in the latter case, memories associated with low levels of reactivation are presumed to be sensitive to other forgetting mechanisms, such as memory decay (Langille, 2019; Frankland et al., 2013).

Although memory reactivation during sleep is generally thought to support systems consolidation (Marr, 1971; McClelland et al., 1995), this computational account (Norman et al., 2006, 2007) would imply that memory reactivation can lead to forgetting when two or more overlapping memories differ in their associative strength. Thus, when two or more memories 'compete' with each other, greater amounts of forgetting will be seen for memories reactivated during sleep, relative to those that are not. An interesting avenue for future research would be to assess the effect of TMR during sleep on the retention of closed-loops. As the retrieval of closed-loops is thought to be more symmetrical, relative to open-loops (Horner et al., 2015; Horner & Burgess, 2014), memory reactivation during sleep should be less likely to lead to retrieval-induced forgetting within a specific closed-loop, as closed-loop elements should be more likely to be reinstated at a similar level of activity. This would imply that memory reactivation of closed-loops during sleep should be relatively holistic and would potentially benefit retention for all associations of a closed-loop.

However, it is likely that different closed-loops all compete for ‘survival’. It is possible that closed-loops are associated with differences in mnemonic strength, such that some closed-loops may be more likely to be reactivated during sleep, relative to others, in the absence of TMR. Memories that are less likely to be reactivated during sleep are thought to be more sensitive to memory decay in the hippocampus (Langille, 2019; Frankland et al., 2013) as decreases in the frequency of reactivation will minimise the likelihood of a memory representation being consolidated in the neocortex (McClelland et al., 1995). As such, differences in the frequency of reactivation during sleep may lead to a pattern of forgetting for closed-loops that tends to be all-or-none; under the assumption that the reactivation of closed-loops during sleep is indeed associated with the concurrent reactivation of all closed-loop elements. Further work is needed to assess this possibility.

5.3.2. Regulatory mechanisms during sleep

The findings presented in Chapter 3 are consistent with computational accounts suggesting that retrieval-induced forgetting can be observed in the absence of any conscious inhibitory control (Norman et al., 2006, 2007). However, further analysis is needed to assess whether oscillatory activity induced by memory reactivation may be related to the effects observed in Chapter 3.

Norman et al., (2006, 2007) argued that activity in the theta band may provide the neural substrate for the selective weakening of competing memories. Consistent with this, Hanslmayr et al. (2010) used a word-stem completion task and observed increased theta power during retrieval practice in a ‘competitive’ (e.g., Fruit: App__ for the exemplar *Apple*), relative to a ‘non-competitive’ condition. (e.g., Fru__: Apple), that subsequently predicted forgetting for non-practice exemplars (e.g., *Banana*). Interestingly, Staudigl, Hanslmayr, and Bäuml (2010) observed that the selective retrieval of a ‘target’ memory was associated with decreases in theta power from the first to the second (and final) retrieval practice trials. This is in line with evidence presented by Kuhl et al. (2007) who demonstrated that BOLD activity in the prefrontal cortex decreased between the first and final retrieval practice trials of an A-B association. Given the role of the prefrontal cortex in inhibitory control, the findings reported by Kuhl et al., (2007) were taken to suggest that decreases in prefrontal activity reflected reductions in interference as a ‘competing’ A-C association was forgotten.

Hanslmayr et al., (2010) also presented evidence for increases in the beta band during retrieval practice, but this increases in beta power did not predict retrieval-induced forgetting. However, Waldhauser et al. (2012) demonstrated that oscillatory activity in the beta band may also be associated with retrieval-induced forgetting. In this study, participants learned a series of shape-colour pairs, where the shape was presented at the centre of fixation and the colour to either the right or left of fixation. During learning, participants would be presented with the shape cue twice, but with the colour presented in the opposite side of fixation each time. Waldhauser et al. (2012) manipulated the amount of interference, such that a ‘new’ colour associate would either be presented in a different (i.e., non-interference condition) or the same colour (i.e., interference condition) during the second presentation of the shape cue. During retrieval practice, participants would be presented with the shape cue and a blank box at either the right or left of fixation. The blank box served as a cue to indicate which colour associate the participants were required to retrieve. Waldhauser et al. (2012) observed increases in beta activity in the cortical hemisphere associated with the opposite visual field as the retrieval cue, but only when the ‘new’ colour associate had been presented in the same colour as the initial shape-colour association. This was taken to suggest that increases in beta activity were associated with the ‘suppression’ of sensory representations of competing shape-colour associations. Importantly, increases in beta power correlated with forgetting in a subsequent retrieval test. This finding is consistent with evidence reported by Wimber et al. (2015), who used fMRI to demonstrate that retrieval practice of an A-B association led to the ‘suppression’ of cortical activity patterns associated with a competing A-C association.

Further analysis of the experiment presented in Chapter 3 will assess the possibility that oscillatory activity in the theta and beta band following memory reactivation correlates with the TMR-induced forgetting effect observed in Chapter 3.

5.3.3. Summary

In Chapter 3, I provided evidence that TMR during sleep can lead to a pattern of forgetting that is (potentially) similar to retrieval-induced forgetting effects seen during wakefulness (see Anderson, 2003 for review). While open-loops appear to be associated with a relatively asymmetrical pattern of forgetting, it is possible that the effect of TMR during sleep for more coherent closed-loops could be associated with a holistic form of forgetting, such that closed-loops that are reinstated during sleep are more likely to be retained in their entirety. In contrast, closed-loops that are not subject to memory reactivation during sleep may be forgotten in a holistic manner, as the absence of memory reactivation will minimise the likelihood of a closed-loop being consolidated in the neocortex, and in turn leave it more susceptible to memory decay in the hippocampus. These proposals can be examined in future experiments assessing the effect of TMR for closed-loops.

5.4. Retrieval following a period of forgetting and consolidation

In Chapter 4, I demonstrated that the retrieval of closed-loops, following a period of forgetting and consolidation, is associated with neocortical activity for all event-elements, including those incidental to the specific retrieval task. However, in contrast to closed-loops retrieved immediately after encoding, the reinstatement of elements incidental to retrieval did not significantly correlate with activity in the hippocampus for closed-loops retrieved after a period of forgetting and consolidation.

5.4.1. Delay-dependent role of the hippocampus

The observation that closed-loops continue to be reinstated holistically - despite there being no evidence for a significant correlation between neocortical reinstatement and activity in the hippocampus - is consistent with the proposal that the forgetting of an initially hippocampal-dependent memory may reflect an interaction between memory decay in the hippocampus and memory consolidation in the neocortex (Frankland et al., 2013; Hardt et al., 2013). Thus, if a neocortical representation of a closed-loop is sufficiently strengthened during systems consolidation, then this closed-loop can perhaps still be retrieved in its entirety; even in instances where a hippocampal representation is subject to decay. As such, systems consolidation, possibly during sleep, may counteract the effects of decay in the hippocampus, allowing for holistic retrieval in the absence of the hippocampus.

Standard consolidation theories (Marr, 1971; McClelland et al., 1995; Squire & Alvarez, 1995) have proposed that memories can become independent of the hippocampus following a period of consolidation. Although some neuroimaging studies have provided support for this proposal (e.g., Niki & Luo, 2002; Piefke et al., 2003; Takashima et al., 2006), it has been argued that these findings are confounded by the fact that recent, relative to remote memories tend to be associated with greater ‘richness’ and/or ‘vividness’ at retrieval (Nadel & Moscovitch, 1997). As such, differences in hippocampal activity for remote and recent memories have been argued to reflect the richness or vividness of retrieval, rather than the age of the memory (Winocur & Moscovitch, 2011). Consistent with this proposal, evidence has demonstrated that hippocampal activity does not consistently differ between the retrieval of recent and remote memories (Addis et al., 2004; Bonnici et al., 2012; Maguire et al., 2001; Ryan et al., 2001), but can instead reflect differences in the vividness of retrieval, irrespective of the remoteness of the memory (Gilboa et al., 2004; Viskontas, Carr, Engel, & Knowlton, 2009). This is consistent with the proposal that memories retrieved independently of the hippocampus may reflect a more gist-like version of the original hippocampal representation (Winocur & Moscovitch, 2011; see also Nadel & Moscovitch, 1997).

In the paradigm used in Chapter 4, participants were explicitly instructed to imagine each of the pairwise associations interacting in some vivid and meaningful way at encoding. I have noted that it is possible that the absence of a significant correlation between the hippocampal activity and the cortical reinstatement of elements incidental to retrieval may reflect a ‘loss’ of the elaborative context established at encoding. Participants may still be able to retrieve closed-loops in a coherent and holistic fashion after a period of consolidation and forgetting, but may perhaps be relying more heavily on gist-like representations in the neocortex. Further work is needed to assess this possibility. However, it has been argued that individuals typically rely on such gist-like representations during retrieval (Brainerd & Reyna, 2002), and it is possible that the likelihood of this increases following a period of forgetting and consolidation.

Although the experiment presented in Chapter 4 does not afford an opportunity to examine this possibility, this would be an intriguing prediction to assess in future work. Future work could incorporate MVPA to assess whether activity patterns reinstated at retrieval become more dissimilar from those seen at encoding, as the interval between encoding and retrieval is increased. Interestingly, Ritchey, Wing, LaBar, and Cabeza (2013) showed that hippocampal activity at retrieval predicted the relationship between memory performance and representational similarity between activity patterns at encoding and immediate retrieval. Future work could assess whether the role of the hippocampus at retrieval decreases as a function of reductions in representational similarity, rather than the interval between encoding and retrieval *per se*.

5.4.2. Conceptual congruence and retention

I have suggested that the forgetting of closed-loops tends to be all-or-none in nature, such that closed-loops are more likely to either be retained or forgotten in their entirety. However, it is not entirely clear why some closed-loops should be more likely to be retained while others are forgotten.

The extent to which new learning is consistent with prior knowledge has been shown to support retention (see Van Kesteren, Ruiter, Fernández, & Henson, 2012 for review) and is thought to accelerate the rate of memory consolidation (Kumaran et al., 2016). It is possible that differences in memory for closed-loops are, at least in part, reflected in their congruence with participants' pre-existing knowledge or the conceptual congruence between the elements themselves. It has long been assumed that congruence increases the 'depth' of encoding, which in turn increases retention (Craik & Tulving, 1975). Consistent with this, Staresina, Gray, & Davachi (2009) showed that congruent word-colour pairs (e.g., *balloon-red*), relative to incongruent pairs (e.g., *elephant-red*) were remembered better on a subsequent memory test. In the experiments presented here, participants were encouraged to imagine the elements interacting as vividly as possible. Although each of the person-location-triplets were created randomly for each participant, some combinations of elements may have been more conceptually congruent than others. It is possible that differences in congruence may decrease or increase participants' ability to semantically elaborate on the interaction between the elements. This may have led to differences in memory for closed-loops. Thus, the pattern of forgetting observed for closed-loops may reflect an interaction between the associative structure formed at encoding and the congruence between the elements that make up a closed-loop. Further work is needed to assess this possibility.

However, congruence with pre-existing knowledge has also been associated with false memory (Spalding, Jones, Duff, Tranel, & Warren, 2015). This is an important point to highlight with reference to the findings in Chapter 4. The ‘non-target’ reinstatement effect, observed in Chapter 4, reflects increases in activity in regions associated with the same category as a non-target element, rather than the reinstatement of the event-specific non-target element. Although the reinstatement effect should be more likely to reflect the retrieval of the associative non-target element (as this is what was learnt at encoding), the likelihood of erroneously reinstating an element from a different event from the same category as the non-target element could be increased if a retrieval cue is strongly congruent with an element from a different event, or if one of the retrieval foils has a strong pre-existing association with an element from the same category as the non-target element. Thus, it is possible that the non-target reinstatement effect observed in Chapter 4 reflects the reinstatement of any element from the same category as the non-target element, rather than the successful reinstatement of the event-specific non-target element.

5.4.3. Summary

In Chapter 4, I provided neuroimaging evidence to show that closed-loops continue to be retrieved in a coherent, and holistic manner following a period of forgetting and consolidation. However, I did not observe evidence for a significant correlation between reinstatement of elements incidental to retrieval and hippocampal activity for closed-loops retrieved following a 24-hour delay. Although this provides some support for standard theories of systems consolidation (McClelland et al., 1995; Squire & Alvarez, 1995), here I have suggested that the absence of a significant correlation at delay may also be related to the ‘vividness’ or ‘richness’ of retrieval. However, this proposal is premature given the evidence presented in Chapter 4, and as such further work is needed to assess this possibility. Further work is also needed to reveal what underlying mechanisms may determine which closed-loops are retained after a period of forgetting. I have raised the possibility that differences in retention may reflect an interaction between the associative structure formed at encoding and the conceptual congruence of the closed-loop itself.

5.5. Conclusion

This thesis provided evidence that the associative structure formed at encoding has a lasting impact on the coherence at retrieval. In relation to closed-loops, the evidence presented here suggests that closed-loops tend to be forgotten in an all-or-none manner, such that closed-loops are more likely to either be retained or forgotten in their entirety. This finding appears most consistent with the proposal that hippocampal representations are forgotten via a process of decay (Frankland et al., 2013; Hardt et al., 2013), and provides evidence to suggest that the rate of decay within a closed-loop is relatively uniform. In contrast, open-loops were shown to be associated with a more asymmetrical pattern of forgetting, similar to retrieval-induced forgetting effects seen during wakefulness (see Anderson, 2003 for review); such that the second encoded associations were forgotten more rapidly as a function of TMR, relative to a non-TMR condition. It is possible that closed-loops show a different pattern of forgetting as a function of TMR, as closed-loops appear to be retrieved in a more coherent manner, relative to open-loops (Horner et al., 2015; Horner & Burgess, 2014). Further work is needed to assess this possibility.

However, consistent with the proposal that closed-loops are forgotten in an all-or-none manner, this thesis presented fMRI evidence that closed-loops retain their coherence at retrieval following a period of forgetting and consolidation. However, the evidence presented here also suggested that the role of the hippocampus in the holistic retrieval of closed-loops may decrease over a period of forgetting and consolidation. This is consistent with standard theories of consolidation suggesting that the role of the hippocampus in retrieval decreases with consolidation (McClelland et al., 1995; Squire & Alvarez, 1995). However, different theories of systems consolidation suggest that event-based representations retrieved independently of the hippocampus differ from those retrieved via the hippocampus (Nadel & Moscovitch, 1997; Winocur & Moscovitch, 2011). It is possible that although closed-loops continue to be retrieved in a coherent manner over time, closed-loops retrieved immediately after encoding and following a period of forgetting and consolidation may differ in some qualitative manner. Further work is needed to assess this possibility.

References

Abel, M., & Bäuml, K. H. T. (2014). Sleep can reduce proactive interference. *Memory*, *22*(4), 332–339.

Achermann, P., & Borbély, A. A. (1997). Low-frequency (< 1 hz) oscillations in the human sleep electroencephalogram. *Neuroscience*, *81*(1), 213–222.

Achim, A. M., & Lepage, M. (2005). Dorsolateral prefrontal cortex involvement in memory post-retrieval monitoring revealed in both item and associative recognition tests. *NeuroImage*, *24*, 1113–1121.

Addis, D. R. (2018). Are episodic memories special? On the sameness of remembered and imagined event simulation. *Journal of the Royal Society of New Zealand*, *48*(2–3), 64–88.

Addis, D. R., Moscovitch, M., Crawley, A. P., & McAndrews, M. P. (2004). Recollective qualities modulate hippocampal activation during autobiographical memory retrieval. *Hippocampus*, *14*, 752–762.

Aggleton, J. P., & Brown, M. W. (1999). Episodic memory, amnesia, and the hippocampal-anterior thalamic axis. *Behavioral and Brain Sciences*, *22*(3), 425–444.

Alvarez, P., & Squire, L. R. (1994). Memory consolidation and the medial temporal lobe: A simple network model. *Proceedings of the National Academy of Sciences of the United States of America*, *91*, 7041–7045.

Anderson, M. C. (2003). Rethinking interference theory: Executive control and the mechanisms of forgetting. *Journal of Memory and Language*, *49*, 415–445.

Anderson, M. C., Bjork, E. L., & Bjork, R. A. (2000). Retrieval-induced forgetting: Evidence for a recall-specific mechanism. *Psychonomic Bulletin and Review*, *7*(3), 522–530.

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., & Spellman, B. A. (1995). On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review*, *102*(1), 68–100.

Antony, J. W., Cheng, L. Y., Brooks, P. P., Paller, K. A., & Norman, K. A. (2018). Competitive learning modulates memory consolidation during sleep. *Neurobiology of Learning and Memory*, *115*, 216–230.

Asch, S. E., & Ebenholtz, S. M. (1962). The Principle of Associative Symmetry. *Proceedings of the American Philosophical Society*, *106*(2), 135–163.

Ashburner, J. (2007). A fast diffeomorphic image registration algorithm. *NeuroImage*, *38*, 95–113.

Ashton, J. E., Cairney, S. A., & Gaskell, M. G. (2018). No effect of targeted memory reactivation during slow-wave sleep on emotional recognition memory. *Journal of Sleep Research*, *27*, 129–137.

Awipi, T., & Davachi, L. (2008). Content-Specific Source Encoding in the Human Medial Temporal Lobe. *Journal of Experimental Psychology: Learning Memory and Cognition*, *34*(4), 769–779.

Axmacher, N., Elger, C. E., & Fell, J. (2008). Ripples in the medial temporal lobe are relevant for human memory consolidation. *Brain*, *131*, 1806–1817.

Backus, A. R., Bosch, S. E., Ekman, M., Grabovetsky, A. V., & Doeller, C. F. (2016). Mnemonic convergence in the human hippocampus. *Nature Communications*, *7*, 11991.

Bakker, A., Kirwan, C. B., Miller, M., & Stark, C. E. L. (2008). Pattern Separation in the Human Hippocampal CA3 and Dentate Gyrus. *Science*, *319*(5870), 1640–1642.

Balaban, H., Assaf, D., Arad Meir, M., & Luria, R. (2019). Different features of real-world objects are represented in a dependent manner in long-term memory. *Journal of Experimental Psychology: General*. Advance online publication

Banino, A., Koster, R., Hassabis, D., & Kumaran, D. (2016). Retrieval-Based Model Accounts for Striking Profile of Episodic Memory and Generalization. *Scientific Reports*, *6*(31330).

Barense, M D, Bussey, T., Lee, A. C. H., Rogers, T. T., Davies, R. R., Saksida, L. M., … Graham, K. S. (2005). Functional specialization in the human medial temporal lobe. *J Neurosci*, *25*(44), 10239–10246.

Barense, M. D., Gaffan, D., & Graham, K. S. (2007). The human medial temporal lobe processes online representations of complex objects. *Neuropsychologia*, *45*(13), 2963–2974.

Barnes, C. A., McNaughton, B. L., Mizumori, S. J. Y., Leonard, B. W., & Lin, L.-H. (1990). Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Progress in Brain Research*, *83*, 287–300.

Barrett, T. R., & Ekstrand, B. R. (1972). Effect of sleep on memory: III. Controlling for time-of-day effects. *Journal of Experimental Psychology*, *96*(2), 321–327.

Barry, D. N., & Maguire, E. A. (2019). Remote Memory and the Hippocampus: A Constructive Critique. *Trends in Cognitive Sciences*, *23*(2), 128–142.

Bäuml, K. H. T., Holterman C., C., & Abel M., M. (2014). Sleep can reduce the testing effect: It enhances recall of restudied items but can leave recall of retrieved items unaffected. *Journal of Experimental Psychology: Learning Memory and Cognition*, *40*(6), 1568–1581.

Ben-Yakov, A., & Dudai, Y. (2011). Constructing realistic engrams: Poststimulus activity of hippocampus and dorsal striatum predicts subsequent episodic memory. *Journal of Neuroscience*, *31*(24), 9032–9042.

Bergmann, T. O., Mölle, M., Diedrichs, J., Born, J., & Siebner, H. R. (2012). Sleep spindle-related reactivation of category-specific cortical regions after learning face-scene associations. *NeuroImage*, *59*, 2733–2742.

Berron, D., Schutze, H., Maass, A., Cardenas-Blanco, A., Kuijf, H. J., Kumaran, D., & Duzel, E. (2016). Strong Evidence for Pattern Separation in Human Dentate Gyrus. *Journal of Neuroscience*, *36*(29), 7569–7579.

Bonnici, H. M., Chadwick, M. J., Lutti, A., Hassabis, D., Weiskopf, N., & Maguire, E. A. (2012). Detecting representations of recent and remote autobiographical memories in vmPFC and hippocampus. *Journal of Neuroscience*, *32*(47), 16982–16991.

Born, J. (2010). Slow-wave sleep and the consolidation of long-term memory. *World Journal of Biological Psychiatry*, *11*, 16–21.

Born, J., & Wilhelm, I. (2012). System consolidation of memory during sleep. *Psychological Research*, *76*(2), 192–203.

Boywitt, C. D., & Meiser, T. (2012). Bound context features are integrated at encoding. *Quarterly Journal of Experimental Psychology*, *65*(8), 1484–1501.

Boywitt, C. D., & Meiser, T. (2012). The role of attention for context-context binding of intrinsic and extrinsic features. *Journal of Experimental Psychology: Learning Memory and Cognition*, *38*(4), 1099–1107.

Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2013). Real-world objects are not represented as bound units: Independent forgetting of different object details from visual memory. *Journal of Experimental Psychology: General*, *142*(3), 791–808.

Brainerd, C. J., & Reyna, V. F. (2002). Fuzzy-trace theory and false memory. *Current Directions in Psychological Science*, *11*(5), 164–169.

Bramão, I., & Johansson, M. (2018). Neural pattern classification tracks transfer-appropriate processing in episodic memory. *ENeuro*, *5*(4), 1–18.

Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network: Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*, *1124*, 1–38.

Buckner, R. L., Koutstaal, W., Schacter, D. L., Wagner, A. D., & Rosen, B. R. (1998). Functional-anatomic study of episodic retrieval using fMRI: I. Retrieval effort versus retrieval success. *NeuroImage*, *7*, 151–162.

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

Buysse, D. J., Reynolds, C. F., Monk, T. H., Berman, S. R., & Kupfer, D. J. (1989). The Pittsburgh sleep quality index: A new instrument for psychiatric practice and research. *Psychiatry Research*, *28*(2), 193–213.

Buzsáki, G. (1996). The hippocampo-neocortical dialogue. *Cerebral Cortex*, *6*, 81–92.

Cairney, S. A., Guttesen, A. á. V., El Marj, N., & Staresina, B. P. (2018). Memory Consolidation Is Linked to Spindle-Mediated Information Processing during Sleep. *Current Biology*, *28*, 948–954.

Cairney, S. A., Lindsay, S., Sobczak, J. M., Paller, K. A., & Gaskell, M. G. (2016). The Benefits of Targeted Memory Reactivation for Consolidation in Sleep are Contingent on Memory Accuracy and Direct Cue-Memory Associations. *Sleep*, *39*(5), 1139–1150.

Cansino, S., Maquet, P., Dolan, R. J., & Rugg, M. D. (2002). Brain Activity Underlying Encoding and Retrieval of Source Memory. *Cerebral Cortex*, *12*(10), 1048–1056.

Cellini, N., & Capuozzo, A. (2018). Shaping memory consolidation via targeted memory reactivation during sleep. *Annals of the New York Academy of Sciences*, *1426*, 52–71.

Cermak, L. S., & O'Connor, M. (1983). The anterograde and retrograde retrieval ability of a patient with amnesia due to encephalitis. *Neuropsychologia*, *21*(3), 213–234.

Chanales, A. J. H., Dudukovic, N. M., Richter, F. R., & Kuhl, B. A. (2019). Interference between overlapping memories is predicted by neural states during learning. *Nature Communications*, *10*(1), 1–12.

Cipolotti, L., Shallice, T., Chan, D., Fox, N., Scahill, R., Harrison, G., … Rudge, P. (2001). Long-term retrograde amnesia... the crucial role of the hippocampus. *Neuropsychologia*, *39*, 151–172.

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

Cohen, N. J., Ryan, J., Hunt, C., Romine, L., Wszalek, T., & Nash, C. (1999). Hippocampal system and declarative (Relational) memory: Summarizing the data from functional neuroimaging studies. *Hippocampus*, *9*(1), 83–98.

Collin, S. H. P., Milivojevic, B., & Doeller, C. F. (2015). Memory hierarchies map onto the hippocampal long axis in humans. *Nature Neuroscience*, *18*(11), 1562–1566.

Conway, M. A. (2005). Memory and the self. *Journal of Memory and Language*, *53*(4), 594–628.

Cooper, R. A., & Ritchey, M. (2019). Cortico-hippocampal network connections support the multidimensional quality of episodic memory. *ELife*, *8*, e45591.

Craik, F. I., & Tulving, E. (1975). Depth of processing and the retention of words in episodic memory. *Journal of Experimental Psychology: General*, *104*(3), 268-294.

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.

Damasio, A. R. (1989). The Brain Binds Entities and Events by Multiregional Activation from Convergence Zones. *Neural Computation*, *1*(1), 123–132.

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.

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

Davachi, L. (2006). Item, context and relational episodic encoding in humans. *Current Opinion in Neurobiology*, *16*(6), 693–700.

Detre, G. J., Natarajan, A., Gershman, S. J., & Norman, K. A. (2013). Moderate levels of activation lead to forgetting in the think/no-think paradigm. *Neuropsychologia*, *51*(12), 2371–2388.

Deuker, L., Olligs, J., Fell, J., Kranz, T. A., Mormann, F., Montag, C., … Axmacher, N. (2013). Memory consolidation by replay of stimulus-specific neural activity. *Journal of Neuroscience*, *33*(49), 19373–19383.

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

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

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

Dobbins, I. G., Foley, H., Schacter, D. L., & Wagner, A. D. (2002). Executive control during episodic retrieval: Multiple prefrontal processes subserve source memory. *Neuron*, *35*, 989–996.

Donaldson, W. (1996). The role of decision processes in remembering and knowing. *Memory and Cognition*, *24*(4), 523–533.

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*, *6*(9), 2591–2599.

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

Dudai, Y., Karni, A., & Born, J. (2015). The Consolidation and Transformation of Memory. *Neuron*, *88*(1), 20–32.

Dupret, D., O'Neill, J., Pleydell-Bouverie, B., & Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nature Neuroscience*, *13*(8), 995–1002.

Düzel, E., Neufang, M., & Heinze, H. J. (2005). The oscillatory dynamics of recognition memory and its relationship to event-related responses. *Cerebral Cortex*, *15*, 1992–2002.

Eichenbaum, H., Yonelinas, A. P., & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annual Review of Neuroscience*, *30*, 123–152.

Eldridge, L. L., Knowlton, B. J., Furmanski, C. S., Bookheimer, S. Y., & Engel, S. A. (2000). Remembering episodes: A selective role for the hippocampus during retrieval. *Nature Neuroscience*, *3*(11), 1149–1152.

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

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. h

Faul, F., Erdfelder, E., Buchner, A., & Lang, A.-G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, *41*(4), 1149–1160.

Fenker, D. B., Schott, B. H., Richardson-Klavehn, A., Heinze, H. J., & Düzel, E. (2005). Recapitulating emotional context: Activity of amygdala, hippocampus and fusiform cortex during recollection and familiarity. *European Journal of Neuroscience*, *21*, 1993–1999.

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

Frankland, P. W., Kohler, S., & Josselyn, S. A. (2013). Hippocampal neurogenesis and forgetting. *Trends in Neurosciences*, *36*(9), 119–130.

Frankland, P. W., & Bontempi, B. (2005). The organization of recent and remote memories. *Nature Reviews Neuroscience*, *6*(2), 119–130.

Gais, S., Albouy, G., Boly, M., Dang-Vu, T. T., Darsaud, A., Desseilles, M., … Peigneux, P.

(2007). Sleep transforms the cerebral trace of declarative memories. *Proceedings of the National Academy of Sciences of the United States of America*, *104*(47), 18778–18783.

Gais, S., & Born, J. (2004). Low acetylcholine during slow-wave sleep is critical for declarative memory consolidation. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(7), 2140–2144.

Gais, S., Lucas, B., & Born, J. (2006). Sleep after learning aids memory recall. *Learning and Memory*, *13*(3), 259–262.

Gais, S., Mölle, M., Helms, K., & Born, J. (2002). Learning-dependent increases in sleep spindle density. *Journal of Neuroscience*, *22*(15), 6830–6834.

Gardiner, J. M., & Java, R. I. (1991). Forgetting in recognition memory with and without recollective experience. *Memory & Cognition*, *19*, 617–623.

Gardner-Medwin, A. R. (1976). The recall of events through the learning of associations between their parts. *Proceedings of the Royal Society of London. Series B - Biological Sciences*, *194*, 375–402.

Geib, B. R., Stanley, M. L., Wing, E. A., Laurienti, P. J., & Cabeza, R. (2017). Hippocampal contributions to the large-scale episodic memory network predict vivid visual memories. *Cerebral Cortex*, *27*(1), 680–693.

Gilboa, A., Winocur, G., Grady, C. L., Hevenor, S. J., & Moscovitch, M. (2004). Remembering our past: Functional neuroanatomy of recollection of recent and very remote personal events. *Cerebral Cortex*, *14*, 1214–1225.

Glickman, S. E. (1961). Perseverative neural processes and consolidation of the memory trace. *Psychological Bulletin*, *58*(3), 218–233.

Göldi, M., van Poppel, E. A. M., Rasch, B., & Schreiner, T. (2019). Increased neuronal signatures of targeted memory reactivation during slow-wave up states. *Scientific*

Reports, *9*, 2715.

Gordon, A. M., Rissman, J., Kiani, R., & Wagner, A. D. (2014). Cortical reinstatement mediates the relationship between content-specific encoding activity and subsequent recollection decisions. *Cerebral Cortex*, *24*, 3350–3364.

Grande, X., Berron, D., Horner, A. J., Bisby, J. A., Düzel, E., & Burgess, N. (2019). Holistic Recollection via Pattern Completion Involves Hippocampal Subfield CA3. *The Journal of Neuroscience*, *39*(41), 8100–8111.

Greeno, J. G., James, C. T., & DaPolito, F. J. (1971). A cognitive interpretation of negative transfer and forgetting of paired associates. *Journal of Verbal Learning and Verbal Behavior*, *10*, 331–345.

Hannula, D. E., Tranel, D., & Cohen, N. J. (2006). The long and the short of it: Relational memory impairments in amnesia, even at short lags. *Journal of Neuroscience*, *26*(32), 8352–8359.

Hanslmayr, S., Staudigl, T., Aslan, A., & Bäuml, K. H. (2010). Theta oscillations predict the detrimental effects of memory retrieval. *Cognitive, Affective and Behavioral Neuroscience*, *10*(3), 329–338.

Hardt, O., Nader, K., & Nadel, L. (2013). Decay happens: The role of active forgetting in memory. *Trends in Cognitive Sciences*, *17*(3), 111–120.

Harlow, I. M., & Donaldson, D. I. (2013). Source accuracy data reveal the thresholded nature of human episodic memory. *Psychonomic Bulletin and Review*, *20*(2), 318–325.

Harlow, I. M., & Yonelinas, A. P. (2016). Distinguishing between the success and precision of recollection. *Memory*, *24*(1), 114–127.

Hassabis, D., & Maguire, E. A. (2007). Deconstructing episodic memory with construction. *Trends in Cognitive Sciences*, *11*(7), 299–306.

Hasselmo, M. E. (1999). Neuromodulation: Acetylcholine and memory consolidation. *Trends in Cognitive Sciences*, *3*(9), 351–359.

Hicks, J. L., & Starns, J. J. (2015). Using Multidimensional Encoding and Retrieval Contexts to Enhance Our Understanding of Stochastic Dependence in Source Memory. *Psychology of Learning and Motivation*, *62*, 101–140.

Hockley, W. E., & Consoli, A. (1999). Familiarity and recollection in item and associative recognition. *Memory and Cognition*, *27*, 657–664.

Hoddes, E., Zarcone, V., Smythe, H., Phillips, R., & Dement, W. C. (1973). Quantification of Sleepiness: A New Approach. *Psychophysiology*, *10*(4), 431–436.

Holdstock, J. S., Mayes, A. R., Gong, Q. Y., Roberts, N., & Kapur, N. (2005). Item recognition is less impaired than recall and associative recognition in a patient with selective hippocampal damage. *Hippocampus*, *15*, 203–215.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, *79*(8), 2554–2558.

Horner, A. J., Bisby, J. A., Bush, D., Lin, W. J., & Burgess, N. (2015). Evidence for holistic episodic recollection via hippocampal pattern completion. *Nature Communications*, *6*, 7462.

Horner, A. J., & Burgess, N. (2013). The associative structure of memory for multi-element events. *Journal of Experimental Psychology: General*, *142*(4), 1370–1383.

Horner, A. J., & Burgess, N. (2014). Pattern completion in multielement event engrams. *Current Biology*, *24*(9), 988–992.

Horner, A. J., & Doeller, C. F. (2017). Plasticity of hippocampal memories in humans. *Current Opinion in Neurobiology*, *24*(9), 988–992.

Horner, A. J., Gadian, D. G., Fuentemilla, L., Jentschke, S., Vargha-Khadem, F., & Duzel, E. (2012). A Rapid, Hippocampus-Dependent, Item-Memory Signal that Initiates Context Memory in Humans. *Current Biology*, *22*(24), 2369–2374.

Hunsaker, M. R., & Kesner, R. P. (2013). The operation of pattern separation and pattern completion processes associated with different attributes or domains of memory. *Neuroscience and Biobehavioral Reviews*, *37*(1), 36–58.

Iber, C., & Iber, C. (2007). *The AASM manual for the scoring of sleep and associated events: rules, terminology and technical specifications*. Westchester, Illinois: American Academy of Sleep Medicine.

Ingram, K. M., Mickes, L., & Wixted, J. T. (2012). Recollection can be weak and familiarity can be strong. *Journal of Experimental Psychology: Learning Memory and Cognition*, *38*(2), 325–339.

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

Johnson, J. D., & Rugg, M. D. (2007). Recollection and the reinstatement of encoding-related cortical activity. *Cerebral Cortex, 17*(11), 2491–2506.

Jones, G. V. (1976). A fragmentation hypothesis of memory: Cued recall of pictures and of sequential position. *Journal of Experimental Psychology: General*, *105*(3), 277–293.

Kan, I. P., Alexander, M. P., & Verfaellie, M. (2009). Contribution of prior semantic knowledge to new episodic learning in amnesia. *Journal of Cognitive Neuroscience*, *21*(5), 938–944.

Kandel, E. R., Dudai, Y., & Mayford, M. R. (2014). The molecular and systems biology of memory. *Cell*, *157*, 163–186.

Kelley, R., & Wixted, J. T. (2001). On the Nature of Associative Information in Recognition

Memory. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *27*(3), 701–722.

Kim, H. (2010). Dissociating the roles of the default-mode, dorsal, and ventral networks in episodic memory retrieval. *NeuroImage*, *50*, 1648–1657.

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

Kirwan, C. B., & Stark, C. E. L. (2007). Overcoming interference: An fMRI investigation of pattern separation in the medial temporal lobe. *Learning and Memory, 14*, 625–633.

Kornell, N., Bjork, R. A., & Garcia, M. A. (2011). Why tests appear to prevent forgetting: A distribution-based bifurcation model. *Journal of Memory and Language*, *65*(2), 85–97.

Kuhl, B. A., Dudukovic, N. M., Kahn, I., & Wagner, A. D. (2007). Decreased demands on cognitive control reveal the neural processing benefits of forgetting. *Nature Neuroscience*, *10*(7), 908–914.

Kuhl, B. A., Johnson, M. K., & Chun, M. M. (2013). Dissociable neural mechanisms for goal-directed versus incidental memory reactivation. *Journal of Neuroscience*, *33*(41), 16099–16109.

Kuhl, B. A., Kahn, I., Dudukovic, N. M., & Wagner, A. D. (2008). Overcoming suppression in order to remember: Contributions from anterior cingulate and ventrolateral prefrontal cortex. *Cognitive, Affective and Behavioral Neuroscience*, *8*(2), 211–221.

Kuhl, B. A., Shah, A. T., Dubrow, S., & Wagner, A. D. (2010). Resistance to forgetting associated with hippocampus-mediated reactivation during new learning. *Nature Neuroscience*, *13*(4), 502–506.

Kumaran, D., Hassabis, D., & McClelland, J. L. (2016). What Learning Systems do Intelligent Agents Need? Complementary Learning Systems Theory Updated. *Trends* *in Cognitive Sciences*, *20*(7), 512–534.

Kumaran, D., Hassabis, D., Spiers, H. J., Vann, S. D., Vargha-Khadem, F., & Maguire, E. A. (2007). Impaired spatial and non-spatial configural learning in patients with hippocampal pathology. *Neuropsychologia*, *45*, 2699–2711.

Kumaran, D., & McClelland, J. L. (2012). Generalization through the recurrent interaction of episodic memories: A model of the hippocampal system. *Psychological Review*, *119*(3), 573–616.

Lahl, O., Wispel, C., Willigens, B., & Pietrowsky, R. (2008). An ultra short episode of sleep is sufficient to promote declarative memory performance. *Journal of Sleep Research*, *17*(1), 3–10.

Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: A practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, *17*(1), 3–10.

Langille, J. J. (2019). Remembering to forget: A dual role for sleep oscillations in memory consolidation and forgetting. *Frontiers in Cellular Neuroscience*, *13*(71), 1–21.

Lau, H., Tucker, M. A., & Fishbein, W. (2010). Daytime napping: Effects on human direct associative and relational memory. *Neurobiology of Learning and Memory*, *93*(4), 554–560.

Lee, A. C. H., Bussey, T. J., Murray, E. A., Saksida, L. M., Epstein, R. A., Kapur, N., … Graham, K. S. (2005). Perceptual deficits in amnesia: Challenging the medial temporal lobe “mnemonic” view. *Neuropsychologia*, *43*(1), 1–11.

Lenth, R. V. (2006). *Java Applets for Power and Sample Size*.

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

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

Maguire, E. A., Henson, R. N. A., Mummery, C. J., & Frith, C. D. (2001). Activity in prefrontal cortex, not hippocampus, varies parametrically with the increasing remoteness of memories. *NeuroReport*, *12*(3), 441–444.

Maldjian, J. A., Laurienti, P. J., Kraft, R. A., & Burdette, J. H. (2003). An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. *NeuroImage*, *19*, 1233–1239.

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

Marshall, L., Helgadóttir, H., Mölle, M., & Born, J. (2006). Boosting slow oscillations during sleep potentiates memory. *Nature*, *444*, 610–613.

Martin, E. (1971). Verbal learning theory and independent retrieval phenomena. *Psychological Review*, *78*(4), 314–332.

Massimini, M., Huber, R., Ferrarelli, F., Hill, S., & Tononi, G. (2004). The sleep slow oscillation as a traveling wave. *Journal of Neuroscience*, *24*(31), 6862–6870.

Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T., & Macrae. (2007). Wandering Minds: The Default Network and Stimulus-Independent Thought. *Science*, *315*, 393–395.

Mayes, A. R., Montaldi, D., & Migo, E. (2007). Associative memory and the medial temporal lobes. *Trends in Cognitive Sciences*, *11*(3), 126–135.

Mayes, A. R., Holdstock, J. S., Isaac, C. L., Montaldi, D., Grigor, J., Gummer, A., … Norman, K. A. (2004). Associative recognition in a patient with selective hippocampal lesions and relatively normal item recognition. *Hippocampus*, *14*(6), 763–784.

McClelland, J. L., & Goddard, N. H. (1996). Considerations arising from a complementary learning systems perspective on hippocampus and neocortex. *Hippocampus*, *6*, 654–665.

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.

McGaugh, J. L. (1966). Time-dependent processes in memory storage. *Science*, *153*(3742), 1351–1358.

McNaughton, B. L., & Morris, R. G. M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences*, *10*(10), 408–415.

Mednick, S. C., McDevitt, E. A., Walsh, J. K., Wamsley, E., Paulus, M., Kanady, J. C., & Drummond, S. P. A. (2013). The critical role of sleep spindles in hippocampal-dependent memory: A pharmacology study. *Journal of Neuroscience*, *13*(10), 4494–4504.

Meiser, T., & Bröder, A. (2002). Memory for Multidimensional Source Information. *Journal of Experimental Psychology: Learning Memory and Cognition*, *28*(1), 116–137.

Meiser, T., & Sattler, C. (2007). Boundaries of the relation between conscious recollection and source memory for perceptual details. *Consciousness and Cognition*, *16*, 189–210.

Meiser, T., Sattler, C., & Weißer, K. (2008). Binding of Multidimensional Context Information as a Distinctive Characteristic of Remember Judgments. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *34*(1), 32–49.

Miller, E. K., & Cohen, J. D. (2001). An Integrative Theory of Prefrontal Cortex Function.

Annual Review of Neuroscience, *24*, 167–202.

Mölle, M., Marshall, L., Gais, S., & Born, J. (2002). Grouping of spindle activity during slow oscillations in human non-rapid eye movement sleep. *Journal of Neuroscience*, *22*(24), 10941–10947.

Moscovitch, M., Nadel, L., Winocur, G., Gilboa, A., & Rosenbaum, R. S. (2006). The cognitive neuroscience of remote episodic, semantic and spatial memory. *Current Opinion in Neurobiology*, *16*, 179–190.

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

Nakazawa, K., Quirk, M. C., Chitwood, R. A., Watanabe, M., Yeckel, M. F., Sun, L. D., … Tonegawa, S. (2002). Requirement for hippocampal CA3 NMDA receptors in associative memory recall. *Science*, *297*, 211–218.

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

Newman, E. L., & Norman, K. A. (2010). Moderate excitation leads to weakening of perceptual representations. *Cerebral Cortex*, *20*, 2760–2770.

Niki, K., & Luo, J. (2002). An fMRI study on the time-limited role of the medial temporal lobe in long-term topographical autobiographic memory. *Journal of Cognitive Neuroscience*, *14*(3), 500–507.

Nilakantan, A. S., Bridge, D. J., VanHaerents, S., & Voss, J. L. (2018). Distinguishing the precision of spatial recollection from its success: Evidence from healthy aging and unilateral mesial temporal lobe resection. *Neuropsychologia*, *119*, 101–106.

Nir, Y., Staba, R. J., Andrillon, T., Vyazovskiy, V. V., Cirelli, C., Fried, I., & Tononi, G.

(2011). Regional Slow Waves and Spindles in Human Sleep. *Neuron*, *70*(1), 153–169.

Norman, K. A., Newman, E., Detre, G., & Polyn, S. (2006). How inhibitory oscillations can train neural networks and punish competitors. *Neural Computation*, *18*, 1577–1610.

Norman, K. A., Newman, E. L., & Detre, G. (2007). A Neural Network Model of Retrieval-Induced Forgetting. *Psychological Review*, *114*(4), 887–953.

Norman, K. A., & O’Reilly, R. C. (2003). Modeling Hippocampal and Neocortical Contributions to Recognition Memory: A Complementary-Learning-Systems Approach. *Psychological Review*, *110*(4), 611–646.

O’Kane, G., Insler, R. Z., & Wagner, A. D. (2005). Conceptual and perceptual novelty effects in human medial temporal cortex. *Hippocampus, 15,* 326–322.

O’Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus, 4*(6), 661–682.

Olman, C. A., Davachi, L., & Inati, S. (2009). Distortion and signal loss in medial temporal lobe. *PLoS ONE, 4*(12), e8160.

Onyper, S. V., Zhang, Y. X., & Howard, M. W. (2010). Some-or-None Recollection: Evidence From Item and Source Memory. *Journal of Experimental Psychology: General, 139*(2), 341–364.

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.

Oyarzún, J. P., Morís, J., Luque, D., de Diego-Balaguer, R., & Fuentemilla, L. (2017).

Targeted memory reactivation during sleep adaptively promotes the strengthening or weakening of overlapping memories. *Journal of Neuroscience*, *37*(32), 7748–7758.

Paivio, A. (1969). Mental imagery in associative learning and memory. *Psychological Review*, *76*(3), 241–263.

Payne, J. D., Stickgold, R., Swanberg, K., & Kensinger, E. A. (2008). Sleep preferentially enhances memory for emotional components of scenes. *Psychological Science*, *19*(8), 781–788.

Peigneux, P., Laureys, S., Fuchs, S., Collette, F., Perrin, F., Reggers, J., … Maquet, P. (2004). Are spatial memories strengthened in the human hippocampus during slow wave sleep? *Neuron*, *44*, 535–545.

Piefke, M., Weiss, P. H., Zilles, K., Markowitsch, H. J., & Fink, G. R. (2003). Differential remoteness and emotional tone modulate the neural correlates of autobiographical memory. *Brain*, *126*, 650–668.

Pihlajamäki, M., Tanila, H., Hänninen, T., Könoönen, M., Mikkonen, M., Jalkanen, V., … Soininen, H. (2003). Encoding of novel picture pairs activates the perirhinal cortex: An fMRI study. *Hippocampus*, *13*, 67–80.

Pihlajamäki, M., Tanila, H., Könönen, M., Hänninen, T., Hämäläinen, A., Soininen, H., & Aronen, H. J. (2004). Visual presentation of novel objects and new spatial arrangements of objects differentially activates the medial temporal lobe subareas in humans. *European Journal of Neuroscience*, *19*, 1939–1949.

Plihal, W., & Born, J. (1997). Effects of early and late nocturnal sleep on declarative and procedural memory. *Journal of Cognitive Neuroscience*, *9*(4), 534–547.

Polyn, S. M., Natu, V. S., Cohen, J. D., & Norman, K. A. (2005). Category-specific cortical activity precedes retrieval during memory search. *Science*, *310*, 1963–1966.

Poppenk, J., Evensmoen, H. R., Moscovitch, M., & Nadel, L. (2013). Long-axis specialization of the human hippocampus. *Trends in Cognitive Sciences*, *17*(5), 230–240.

Poppenk, J., & Moscovitch, M. (2011). A hippocampal marker of recollection memory ability among healthy young adults: Contributions of posterior and anterior segments. *Neuron*, *72*, 931–937.

Postman, L. (1971). Transfer, interference and forgetting. *Experimental Psychology*, 1019–1132.

Potter, K. W., Huszar, L. D., & Huber, D. E. (2018). Does inhibition cause forgetting after selective retrieval? A reanalysis and failure to replicate. *Cortex*, *104*, 26–45.

Poulton, E. C. (1982). Influential companions: Effects of one strategy on another in the within-subjects designs of cognitive psychology. *Psychological Bulletin*, *91*(3), 673–690.

Preston, A. R., Shrager, Y., Dudukovic, N. M., & Gabrieli, J. D. E. (2004). Hippocampal contribution to the novel use of relational information in declarative memory. *Hippocampus*, *14*, 148–152.

Ranganath, C. (2010). Binding items and contexts: The cognitive neuroscience of episodic memory. *Current Directions in Psychological Science*, *19*(3), 131–137.

Ranganath, C., Cohen, M. X., Dam, C., & D’Esposito, M. (2004). Inferior Temporal, Prefrontal, and Hippocampal Contributions to Visual Working Memory Maintenance and Associative Memory Retrieval. *Journal of Neuroscience*, *24*(16), 3917–3925.

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

Rasch, B., Büchel, C., Gais, S., & Born, J. (2007). Odor cues during slow-wave sleep prompt declarative memory consolidation. *Science*, *315*, 1426–1429.

Rauchs, G., Feyers, D., Landeau, B., Bastin, C., Luxen, A., Maquet, P., & Collette, F. (2011). Sleep contributes to the strengthening of some memories over others, depending on hippocampal activity at learning. *Journal of Neuroscience*, *31*(7), 2563–2568.

Renoult, L., Irish, M., Moscovitch, M., & Rugg, M. D. (2019). From Knowing to Remembering: The Semantic–Episodic Distinction. *Trends in Cognitive Sciences*, *23*(12), 1041-1057.

Reyna, V. F., & Brainerd, C. J. (1995). Fuzzy-trace theory: An interim synthesis. *Learning and Individual Differences*, *7*(1), 1–75.

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*, e18260.

Ritchey, M., Montchal, M. E., Yonelinas, A. P., & Ranganath, C. (2015). Delay-dependent contributions of medial temporal lobe regions to episodic memory retrieval. *ELife*, *4*,

Ritchey, M., Wing, E. A., LaBar, K. S., & Cabeza, R. (2013). Neural similarity between encoding and retrieval is related to memory via hippocampal interactions. *Cerebral Cortex*, *23*(12), 2818–2828.

Roediger, H. L., & Butler, A. C. (2011). The critical role of retrieval practice in long-term retention. *Trends in Cognitive Sciences*, *15*(1), 20–27.

Roediger, H. L., & Karpicke, J. D. (2006). Test-Enhanced Learning. *Psychological Science*, *17*(3), 249–255.

Roy, D. S., Arons, A., Mitchell, T. I., Pignatelli, M., Ryan, T. J., & Tonegawa, S. (2016).

Memory retrieval by activating engram cells in mouse models of early Alzheimer’s disease. *Nature*, *531*(7595), 508–512.

Rudoy, J. D., Voss, J. L., Westerberg, C. E., & Paller, K. A. (2009). Strengthening individual memories by reactivating them during sleep. *Science*, *326*(5956), 1079.

Rugg, M. D., & Vilberg, K. L. (2013). Brain networks underlying episodic memory retrieval. *Current Opinion in Neurobiology*, *23*, 255–260.

Ryan, J. D., Althoff, R. R., Whitlow, S., & Cohen, N. J. (2000). Amnesia is a deficit in relational memory. *Psychological Science*, *11*(6), 454–461.

Ryan, L., Nadel, L., Keil, K., Putnam, K., Schnyer, D., Trouard, T., & Moscovitch, M. (2001). Hippocampal complex and retrieval of recent and very remote autobiographical memories: Evidence from functional magnetic resonance imaging in neurologically intact people. *Hippocampus*, *11*, 707–714.

Ryan, T. J., Roy, D. S., Pignatelli, M., Arons, A., & Tonegawa, S. (2015). Engram cells retain memory under retrograde amnesia. *Science*, *348*(6238), 1007-1013.

Sadeh, T., Ozubko, J. D., Winocur, G., & Moscovitch, M. (2014). How we forget may depend on how we remember. *Trends in Cognitive Sciences*, *18*(1), 26–36.

Sadeh, T., Ozubko, J. D., Winocur, G., & Moscovitch, M. (2016). Forgetting Patterns Differentiate Between Two Forms of Memory Representation. *Psychological Science*, *27*(6), 810–820.

Saletin, J. M., Goldstein, A. N., & Walker, M. P. (2011). The role of sleep in directed forgetting and remembering of human memories. *Cerebral Cortex*, *21*, 2534–2541.

Sanders, H. I., & Warrington, E. K. (1971). Memory for remote events in amnesic patients. *Brain*, *94*, 661–668.

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

Schacter, D. L., Addis, D. R., & Buckner, R. L. (2007). Remembering the past to imagine the future: The prospective brain. *Nature Reviews Neuroscience, 8*(657–661), 657–661.

Schapiro, A. C., McDevitt, E. A., Rogers, T. T., Mednick, S. C., & Norman, K. A. (2018). Human hippocampal replay during rest prioritizes weakly learned information and predicts memory performance. *Nature Communications, 9*(1), 3920.

Schmidt, C., Peigneux, P., Muto, V., Schenkel, M., Knoblauch, V., Münch, M., … Cajochen, C. (2006). Encoding difficulty promotes postlearning changes in sleep spindle activity during napping. *Journal of Neuroscience*, *26*(35), 8976–8982.

Schnider, A., Bassetti, C., Gutbrod, K., & Ozdoba, C. (1995). Very severe amnesia with acute onset after isolated hippocampal damage due to systemic lupus erythematosus. *Journal of Neurology, Neurosurgery and Psychiatry*, *59*(6), 644–646.

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

Schreiner, T., Lehmann, M., & Rasch, B. (2015). Auditory feedback blocks memory benefits of cueing during sleep. *Nature Communications*, *6*, 8729.

Schreiner, T., & Rasch, B. (2015). Boosting vocabulary learning by verbal cueing during sleep. *Cerebral Cortex*, *25*, 4169–4179.

Schuetze, B. A., Eglington, L. G., & Kang, S. H. K. (2019). Retrieval practice benefits memory precision. *Memory*, *27*(8), 1091–1098.

Scoville, W. B., & Milner, B. (1957). Loss of recentmemory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery & Psychiatry, 20*(1), 11–21.

Sekeres, M. J., Bonasia, K., St-Laurent, M., Pishdadian, S., Winocur, G., Grady, C., & Moscovitch, M. (2016). Recovering and preventing loss of detailed memory: Differential rates of forgetting for detail types in episodic memory. *Learning and Memory*, *23*(15), 72–82.

Shohamy, D., & Wagner, A. D. (2008). Integrating Memories in the Human Brain: Hippocampal-Midbrain Encoding of Overlapping Events. *Neuron*, *60*(2), 378–389.

Siapas, A. G., & Wilson, M. A. (1998). Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron*, *21*, 1123–1128.

Simon, K. C. N. S., Gómez, R. L., & Nadel, L. (2018). Losing memories during sleep after targeted memory reactivation. *Neurobiology of Learning and Memory*, *151*, 10–17.

Sirota, A., & Buzsáki, G. (2005). Interaction between neocortical and hippocampal networks via slow oscillations. *Thalamus and Related Systems*, *3*(4), 245–259.

Sirota, A., Csicsvari, J., Buhl, D., & Buzsáki, G. (2003). Communication between neocortex and hippocampus during sleep in rodents. *Proceedings of the National Academy of Sciences of the United States of America*, *100*(4), 2065–2069.

Song, Z., Wixted, J. T., Hopkins, R. O., & Squire, L. R. (2011). Impaired capacity for familiarity after hippocampal damage. *Proceedings of the National Academy of Sciences*, *108*(23), 9655–9660.

Spalding, K. N., Jones, S. H., Duff, M. C., Tranel, D., & Warren, D. E. (2015). Investigating the neural correlates of schemas: Ventromedial prefrontal cortex is necessary for normal schematic influence on memory. *Journal of Neuroscience*, *35*(47), 15746–15751.

Spaniol, J., Davidson, P. S. R., Kim, A. S. N., Han, H., Moscovitch, M., & Grady, C. L. (2009). Event-related fMRI studies of episodic encoding and retrieval: Meta-analyses using activation likelihood estimation. *Neuropsychologia*, *47*, 1765–1779.

Squire, L. R., Haist, F., & Shimamura, A. P. (1989). The neurology of memory: Quantitative assessment of retrograde amnesia in two groups of amnesic patients. *Journal of Neuroscience*, *9*(3), 828–839.

Squire, L. R., & Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science*, *253*(5026), 1380–1386.

Squire, L. R., & Alvarez, P. (1995). Retrograde amnesia and memory consolidation: a neurobiological perspective. *Current Opinion in Neurobiology*, *5*(2), 169–177.

St-Laurent, M., Moscovitch, M., Levine, B., & McAndrews, M. P. (2009). Determinants of autobiographical memory in patients with unilateral temporal lobe epilepsy or excisions. *Neuropsychologia*, *47*, 2211–2221.

Staba, R. J., Wilson, C. L., Bragin, A., & Fried, I. (2002). Quantitative analysis of high-frequency oscillations (80-500 Hz) recorded in human epileptic hippocampus and entorhinal cortex. *Journal of Neurophysiology*, *88*, 1743–1752.

Staresina, B. P., Alink, A., Kriegeskorte, N., & Henson, R. N. (2013). Awake reactivation predicts memory in humans. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(52), 21159–21164.

Staresina, B. P., Bergmann, T. O., Bonnefond, M., Van Der Meij, R., Jensen, O., Deuker, L., … Fell, J. (2015). Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus during sleep. *Nature Neuroscience*, *18*(11), 1679–1686.

Staresina, B. P., Cooper, E., & Henson, R. N. (2013). Reversible information flow across the medial temporal lobe: The hippocampus links cortical modules during memory retrieval. *Journal of Neuroscience*, *33*(35), 14184–14192.

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

Staresina, B. P., & Davachi, L. (2008). Selective and shared contributions of the hippocampus and perirhinal cortex to episodic item and associative encoding. *Journal of Cognitive Neuroscience*, *20*(8), 1478–1489.

Staresina, B. P., Duncan, K. D., & Davachi, L. (2011). Perirhinal and parahippocampal cortices differentially contribute to later recollection of object- and scene-related event details. *Journal of Neuroscience*, *31*(24), 8739–8747.

Staresina, B. P., Gray, J. C., & Davachi, L. (2009). Event congruency enhances episodic memory encoding through semantic elaboration and relational binding. *Cerebral Cortex*, *19*(5), 1198–1207.

Staresina, B. P., Henson, R. N. A., Kriegeskorte, N., & Alink, A. (2012). Episodic reinstatement in the medial temporal lobe. *Journal of Neuroscience*, *32*(50), 18150–18156.

Starns, J. J., & Hicks, J. L. (2005). Source Dimensions Are Retrieved Independently in Multidimensional Monitoring Tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(6), 1213–1220.

Starns, J. J., & Hicks, J. L. (2008). Context attributes in memory are bound to item information, but not to one another. *Psychonomic Bulletin & Review*, *15*(2), 309–314.

Staudigl, T., Hanslmayr, S., & Bäuml, K. H. T. (2010). Theta oscillations reflect the dynamics of interference in episodic memory retrieval. *Journal of Neuroscience*, *30*(34), 11356–11362.

Steinvorth, S., Levine, B., & Corkin, S. (2005). Medial temporal lobe structures are needed to re-experience remote autobiographical memories: Evidence from H.M. and W.R. *Neuropsychologia*, *43,* 479–496.

Sterpenich, V., Albouy, G., Darsaud, A., Schmidt, C., Vandewalle, G., Dang Vu, T. T., … Maquet, P. (2009). Sleep promotes the neural reorganization of remote emotional memory. *Journal of Neuroscience*, *29*(16), 5143–5152.

Stickgold, R., & Walker, M. P. (2013). Sleep-dependent memory triage: evolving generalization through selective processing. *Nature Neuroscience*, *16*(2), 139–145.

Sun, S. Z., Fidalgo, C., Barense, M. D., Lee, A. C. H., Cant, J. S., & Ferber, S. (2017). Erasing and blurring memories: The differential impact of interference on separate aspects of forgetting. *Journal of Experimental Psychology: General*, *146*(11), 1606–1630.

Sutterer, D. W., & Awh, E. (2016). Retrieval practice enhances the accessibility but not the quality of memory. *Psychonomic Bulletin and Review*, *23*(3), 831–841.

Takashima, A., Petersson, K. M., Rutters, F., Tendolkar, I., Jensen, O., Zwarts, M. J., … Fernández, G. (2006). Declarative memory consolidation in humans: A prospective functional magnetic resonance imaging study. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(3), 756–761.

Takashima, A., Nieuwenhuis, I. L. C., Jensen, O., Talamini, L. M., Rijpkema, M., & Fernández, G. (2009). Shift from hippocampal to neocortical centered retrieval network with consolidation. *Journal of Neuroscience*, *29*(32), 10087–10093.

Tanji, K., Suzuki, K., Fujii, T., Higano, S., & Yamadori, A. (2003). A case of frontal network amnesia. *Journal of Neurology Neurosurgery and Psychiatry*, *74*, 106–109.

Taylor, K. I., Moss, H. E., Stamatakis, E. A., & Tyler, L. K. (2006). Binding crossmodal object features in perirhinal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(21), 8239–8244.

Teyler, T. J., & DiScenna, P. (1986). The hippocampal memory indexing theory. *Behavioral*

Neuroscience, *100*(2), 147–154.

Treves, A., & Rolls, E. T. (1992). Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus*, *2*(2), 189–199.

Treves, A., & Rolls, E. T. (1994). Computational analysis of the role of the hippocampus in memory. *Hippocampus*, *4*(3), 374–391.

Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., … Morris, R. G. M. (2007). Schemas and memory consolidation. *Science*, *316*, 76–82.

Tse, D., Takeuchi, T., Kakeyama, M., Kajii, Y., Okuno, H., Tohyama, C., … Morris, R. G. M. (2011). Schema-dependent gene activation and memory encoding in neocortex. *Science*, *333*, 891–895.

Tucker, M., Hirota, Y., Wamsley, E., Lau, H., Chaklader, A., & Fishbein, W. (2006). A daytime nap containing solely non-REM sleep enhances declarative but not procedural memory. *Neurobiology of Learning and Memory*, *86*(2), 241–247.

Tulving, E. (1983). *Elements of Episodic Memory.* Oxford, UK: Clarendon Press.

Tulving, E. (1985). Memory and Consciousness. *Canadian Psychology*, *26*(1), 1–12.

Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, *5*(4), 381–391.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., … Joliot, M. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *NeuroImage*, *15*, 273–289.

Uncapher, M. R., Otten, L. J., & Rugg, M. D. (2006). Episodic Encoding Is More than the Sum of Its Parts: An fMRI Investigation of Multifeatural Contextual Encoding.

Neuron, *52*(3), 547–556.

Uncapher, M. R., & Rugg, M. D. (2005). Encoding and the durability of episodic memory: A functional magnetic resonance imaging study. *Journal of Neuroscience*, *25*(31), 7260–7267.

Underwood, B. J. (1957). Interference and forgetting. *Psychological Review, 64*(1), 49–60.

Vaidya, C. J., Zhao, M., Desmond, J. E., & Gabrieli, J. D. E. (2002). Evidence for cortical encoding specificity in episodic memory: Memory-induced re-activation of picture processing areas. *Neuropsychologia*, *40*, 2136–2143.

van de Ven, G. M., Trouche, S., McNamara, C. G., Allen, K., & Dupret, D. (2016). Hippocampal Offline Reactivation Consolidates Recently Formed Cell Assembly Patterns during Sharp Wave-Ripples. *Neuron*, *92*, 968–974.

Van Dongen, E. V., Takashima, A., Barth, M., Zapp, J., Schad, L. R., Paller, K. A., & Fernández, G. (2012). Memory stabilization with targeted reactivation during human slow-wave sleep. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(26), 10575–10580.

Van Kesteren, M. T. R., Ruiter, D. J., Fernández, G., & Henson, R. N. (2012). How schema and novelty augment memory formation. *Trends in Neurosciences, 35*(4), 211–219.

Vargha-Khadem, F., Gadian, D. G., Watkins, K. E., Connelly, A., Van Paesschen, W., & Mishkin, M. (1997). Differential Effects of Early Hippocampal Pathology on Episodic and Semantic Memory. *Science, 277*(5324), 376–380.

Viskontas, I. V., Carr, V. A., Engel, S. A., & Knowlton, B. J. (2009). The neural correlates of recollection: Hippocampal activation declines as episodic memory fades. *Hippocampus*, *19*, 265–272.

Viskontas, I. V., Knowlton, B. J., Steinmetz, P. N., & Fried, I. (2006). Differences in

Mnemonic Processing by Neurons in the Human Hippocampus and Parahippocampal Regions. *Journal of Cognitive Neuroscience*, *18*(10), 1654–1662.

Wagner, U., Gais, S., Haider, H., Verleger, R., & Born, J. (2004). Sleep inspires insight. *Nature*, *427*(6972), 352–355.

Wais, P. E., Wixted, J. T., Hopkins, R. O., & Squire, L. R. (2006). The Hippocampus Supports both the Recollection and the Familiarity Components of Recognition Memory. *Neuron*, *49*(3), 459–466.

Waldhauser, G. T., Johansson, M., & Hanslmayr, S. (2012). Alpha/beta oscillations indicate inhibition of interfering visual memories. *Journal of Neuroscience*, *32*(6), 1953–1961.

Watanabe, T., Kimura, H. M., Hirose, S., Wada, H., Imai, Y., Machida, T., … Konishi, S. (2012). Functional dissociation between anterior and posterior temporal cortical regions during retrieval of remote memory. *Journal of Neuroscience*, *32*(28), 9659–6970.

Watson, H. C., Wilding, E. L., & Graham, K. S. (2012). A Role of the Perirhinal Cortex in Memory for Novel Object-Context Associations. *Journal of Neuroscience*, *32*(13), 4473–4481.

Weis, S., Specht, K., Klaver, P., Tendolkar, I., Willmes, K., Ruhlmann, J., … Fernandez, G. (2004). Process dissociation between contextual retrieval and item recognition. *NeuroReport*, *15*(18), 2729–2733.

Weiskopf, N., Hutton, C., Josephs, O., & Deichmann, R. (2006). Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: A whole-brain analysis at 3 T and 1.5 T. *NeuroImage*, *33*, 493–504.

Wheeler, M. E., & Buckner, R. L. (2004). Functional-anatomic correlates of remembering and knowing. *NeuroImage*, *21*, 1337–1349.

Wheeler, M. E., Petersen, S. E., & Buckner, R. L. (2000). Memory's echo: Vivid remembering reactivates sensory-specific cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *97*(20), 11125–11129.

Wichawut, C., & Martin, E. (1971). Independence of A-B and A-C associations in retroaction. *Journal of Verbal Learning and Verbal Behavior*, *10*, 316–321.

Wilhelm, I., Diekelmann, S., Molzow, I., Ayoub, A., Mölle, M., & Born, J. (2011). Sleep selectively enhances memory expected to be of future relevance. *Journal of Neuroscience*, *31*(5), 1563–1569.

Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, *265*, 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., Rutschmann, R. M., Greenlee, M. W., & Bäuml, K. H. (2008). Retrieval from episodic memory: Neural mechanisms of interference resolution. *Journal of Cognitive Neuroscience*, *21*(3), 538–549.

Winocur, G., & Moscovitch, M. (2011). Memory transformation and systems consolidation. *Journal of the International Neuropsychological Society*, *17*, 766–780.

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*, 2339–2356.

Wixted, J. T. (2004). The Psychology and Neuroscience of Forgetting. *Annual Review of Psychology*, *55*(1), 235–269.

Wixted, J. T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*, *114*(1), 152–176.

Wixted, J. T., & Squire, L. R. (2004). Recall and recognition are equally impaired in patients with selective hippocampal damage. *Cognitive, Affective and Behavioral Neuroscience*,

Wixted, J. T., & Stretch, V. (2004). In defense of the signal detection interpretation of remember/know judgments. *Psychonomic Bulletin and Review*, *11*(4), 616–641.

Woodruff, C. C., Johnson, J. D., Uncapher, M. R., & Rugg, M. D. (2005). Content-specificity of the neural correlates of recollection. *Neuropsychologia*, *43*, 1022–1032.

Woods, B. T., Schoene, W., & Kneisley, L. (1982). Are hippocampal lesions sufficient to cause lasting amnesia? *Journal of Neurology Neurosurgery and Psychiatry*, *45*, 243–247.

Yassa, M. A., & Stark, C. E. L. (2011). Pattern separation in the hippocampus. *Trends in Neurosciences*, *34*(10), 515–525.

Yonelinas, A. P. (1994). Receiver-operating characteristics in recognition memory: Evidence for a dual-process model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(6), 1341–1354.

Yonelinas, A. P. (2002). The nature of recollection and familiarity: A review of 30 years of research. *Journal of Memory and Language*, *46*, 441–517.

Yonelinas, A. P., Aly, M., Wang, W. C., & Koen, J. D. (2010). Recollection and familiarity: Examining controversial assumptions and new directions. *Hippocampus*, *20*(11), 1178–1194.

Zeithamova, D., Dominick, A. L., & Preston, A. R. (2012). Hippocampal and Ventral Medial Prefrontal Activation during Retrieval-Mediated Learning Supports Novel Inference. *Neuron*, *75*(1), 168–179.