Piecing it together: Examining the processing of single features and feature bindings in visual working memory

David Thomas Goldenhaus-Manning

Thesis submitted for the degree of Doctor of Philosophy (PhD)

Department of Psychology

University of Essex

February 2024

Acknowledgement

I would like to thank my husband Andrew for his continued love and patience during this difficult journey. I would also like to thank my fantastic supervisors Dr Vanessa Loaiza and Dr Nicholas Cooper for their constant support and encouragement. Finally, I would like to thank fellow PhD students Charli Sherman and Hiu Wah Cheung for their support.

Abstract

The goal of this thesis was to examine how visual features are prioritised during goal-orientated processing in working memory. It was examined whether attention is required for the maintenance of features in visual working memory (Chapters 2 and 3), and whether there are differences between the prioritisation of location and non-location features during goal-orientated processing (Chapter 4). Additionally, it was examined whether during the prioritisation of updated information, outdated information is temporarily (i.e., outdated information is still present and accessible) or permanently (i.e., outdated information is no longer present or accessible) removed (Chapters 5 and 6). Chapters 2 and 3 observed that manipulating attention did not affect the likelihood of binding errors occurring (i.e., the probability of a feature becoming misbound). However, it was observed that the likelihood of a target feature being in memory (target memory) was affected by manipulating attention. In Chapter 4 it was observed that there are no distinct differences between the prioritisation of location and non-location features during goal-oriented processing. Similarly to Experiment 2 of Chapter 2, it was observed that guiding attention improved target memory. Chapters 5 and 6 investigated whether the prioritisation of outdated information relies on the temporary or permanent removal of outdated information, by altering the to-be-updated feature type. In both Chapters 5 and 6, it was observed that the model simulating temporary removal best fit the data. Thus, prioritising updated information relies on temporary removal, irrespective of the feature type. Therefore, the findings of this thesis suggest that the processing of visual information occurs in an object-based manner, but additional resources aid the prioritisation of goal-relevant feature bindings.

Contents

List of Figures

List of Tables

Appendices

Chapter 1

How do we keep visual information active in our immediate memory?

What is meant by "working memory" and why is it important?

Research concerned with how human memory operates has been a great interest over the last century, particularly, one's ability to actively hold onto information to achieve a goal (i.e., goal-orientated processing). Early research concerned with how information is kept active suggested that human memory comprises a short-term memory store (where information is temporarily held active) and a long-term store (i.e., relevant information is stored for future use; Atkinson & Shiffrin, 1968). Atkinson and Shiffrin (1968) posited that initially, after information is registered sensorily, it is then encoded into a temporary short-term store or a long-term store for future recall. Arguably Atkinson and Shiffrin's (1968) explanation of human memory was a fundamental steppingstone for understanding the processing of information. However, their explanation of short-term memory was too passive to capture how temporarily stored information may be held active to allow for it to be manipulated to achieve one's goal.

In contrast to short-term memory, working memory (WM) is thought to be a more dynamic limited capacity system that allows for the temporary maintenance and manipulation of information for goal-orientated processing Cowan et al., 2020). Understanding how WM is conceptualised and how it operates is important, especially because it is applied to other aspects of human behaviour. For example, WM capacity (that is, the number of items one can hold active in one's WM) is strongly related to general intelligence (Colom et al., 2005). Furthermore, research has demonstrated that applying WM training strategies in an educational setting improves the processing of verbal information in young children (St Clair-Thompson et al., 2010). This therefore demonstrates the diversity of WM research is not only relevant to memory researchers but is also applicable to the greater understanding of human cognition. Nevertheless, there is an ongoing debate in WM research on a variety of issues, such as whether WM is domain-specific (i.e., information is processed in separate, specific, stores) or domain-general (i.e., information is processed by a holistic memory system rather than in separate stores), and how we process information in WM.

Domain-specific perspectives of WM

Domain-specific perspectives of WM suggest that WM may be organised such that information is processed in specific stores (i.e., visual information is processed in a visual store, and verbal information in a verbal store). According to the most influential domain-specific perspective, Baddeley and Hitch's (1974) multi-component model, information in WM is separately stored and maintained in domain-specific modules: language-based information is stored and maintained in the phonological loop, and visual-spatial information is stored and maintained in the visual-spatial sketchpad. Furthermore, the researchers posited that the processing of verbal and visual information is governed by the central executive (also see Baddeley, 1996). In a discussion of the role of the central executive, Baddeley (1996) argued that the central executive may be responsible for limiting the number of items that can be retained in WM. Baddeley (1996) argued that the central executive is responsible for one's ability to concurrently carry out two tasks (such as in a dual-task paradigm). In this case, it was argued that the central executive acts as a general information processing system that selectively attends to goal-relevant information and discards non-relevant information. This therefore suggests that the central executive may govern the limited capacity of WM by attending to only goal-relevant information during processing.

Moreover, he posited that the central executive acts as a direct link between long-term memory (LTM) and WM. This suggests that WM and LTM are separate stores in the multi-component model. In comparison to Atkinson and Shiffrin’s (1968) model, the multi-component model offers a wider understanding of how information is processed in human memory, with domain-specific stores and a governing component that directs information processing. For example, early research investigating the multi-component model primarily focused on verbal processing in the phonological loop (e.g., Baddeley et al., 1984). In a series of verbal WM experiments, Baddeley and colleagues (1984) demonstrated that the phonological similarity effect (i.e., poorer immediate memory for phonologically similar words) was unaffected by articulatory suppression (i.e., repeating distinct words or phrases out loud). The researchers suggested that verbal information is processed in WM despite the suppression of a verbal store through articulatory suppression. This suggests that the processing of verbal information in the phonological loop may not be limited by distractors of the same domain-type. Baddeley and colleagues suggested that access to the phonological loop is strengthened through auditory presentation of verbal information. Thus it may be the case that the articulatory suppression allowed for goal-relevant verbal information to be stored correctly rather than disrupting its storage.

Further work concerning the visuospatial sketchpad proposed that visual and spatial information is processed in their respective domain-specific sub-systems of the visuo-spatial sketchpad known as the visual cache and the inner scribe (Logie, 1995, 2003; Logie & Pearson, 1997). It was argued that non-location information (i.e., colour or shape) are processed in the (passive) visual cache, whereas location information is processed in the inner scribe (Logie, 1995). Moreover, it has been argued that processing location information through the inner scribe may rely on attention (Logie, 2003). A case study investigating two patients who were unable to mentally represent items were in fact able to mentally rotate a series of real-world objects (i.e., a vase, coffee pot, ashtray, and a bottle) without requiring spatial processing from LTM. This suggested that the processing of visual information may require attentional resources to rotate mentally rotate the objects, but the deficit of keeping information active may be due to a faulty visual cache (Beschin et al., 2005). I will return to this point regarding whether attentional resources are required for processing visual information, particularly location-based information, in later sections.

One way in which advocates of the multi-component model have investigated whether information is processed in separate, domain-specific stores is through dual-task paradigms. In a dual-task paradigm, a participant is required to complete a visual or verbal memory task alongside a second visual or verbal task (usually before recalling an aspect of the first memory task). The participant may be tested on the first task alone, the second task alone, or both tasks (Cowan & Morey, 2007). For example, Logie and Marchetti (1991) observed that during a visual and spatial memory task, disrupting the retention of the to-be-retained visual and spatial information, by means of a secondary visual and spatial task, respectively, specifically increased incorrect mean memory recognition for visual and spatial information. Moreover, the researchers observed that disrupting the to-be-retained visual and spatial information by means of the opposite secondary tasks did not increase incorrect mean memory recognition as the latter. The researchers therefore argued that this indicates that visual information is stored in a visual-specific store and spatial information is stored in a spatial-specific store. This is therefore aligned with Logie's (1995) view that visual and spatial information is stored and maintained in domain-specific stores.

Further evidence from dual-task paradigms has suggested that a participant's ability to concurrently retain visual and verbal information indicates that visual and verbal information is stored in domain-specific stores (Cocchini et al., 2002). For example, in Experiment 1 of Cocchini and colleagues (2002), participants were either presented with a series of digits in a sequential order followed by a visual pattern, or a series of digits and a spatial tracking task. The order of the tasks was randomised, in that, participants may be presented with the verbal task first, or the visual task first. The researchers observed that recalling the digit sequence first was impacted by concurrently storing visual information. Similarly, the researchers observed that recalling the pattern first was impacted by concurrently storing digits. This suggests that visual and verbal information is stored in domain-specific stores but retaining visual and verbal information simultaneously comes at a cost during the retrieval of domain-specific information.

Extending the multi-component model, Baddeley (2000) proposed that a new fourth component (known as the episodic buffer) that is modality-general and can temporarily hold information active for further processing and storage within the respective domain-specific stores. Baddeley (2000) posited that in line with the visuo-spatial sketchpad and the phonological loop, the episodic buffer is governed by the central executive. However, unlike the visuo-spatial sketchpad and the phonological loop, Baddeley (2000) argued that the episodic buffer is not a domain-specific component, rather it is acts as the ‘middleman’ between the processing of visual and verbal information. For example, research has argued that the episodic buffer is able to bind visual and verbal information together across their respective domain-specific stores (Allen et al., 2015; also see Darling et al., 2017). Moreover, it has been argued that the episodic buffer holds goal-relevant verbal and visual information active by its multidimensional coding (see Baddeley, 2012 for a review of the multi-component model). Baddeley and colleagues (2010) further demonstrated that the episodic buffer does not require effortful cognition to engage its ability to keep information active in WM. Moreover, the researchers posited that the episodic buffer may temporarily store information from LTM. In a recent review of the multi-component model, Baddeley and colleagues (2021) argued that the episodic buffer may be an attentional system that is able to selectively maintain goal-relevant information from its respective stores. This suggests that the episodic buffer maintains information irrespective of its modality. As I will discuss in the next section, this draws parallels with domain-general models of WM.

In summary, domain-specific models such as the multi-component model propose that the processing and storage of information depends on the type of information that is to-be-retained (i.e., verbal, or visual). However, a primary challenge for domain-specific views is the great deal of work suggesting that WM is more domain-general than domain-specific in its functioning, as I will explain further in the next section.

Domain-general perspectives of WM

Although the multi-component model was among the first models to discuss the inner workings of active information processing, storage and maintenance, alternative models have hypothesised different approaches towards WM. A second view of WM suggests that WM may be conceptualised as domain-general, wherein information is processed across a flexible resource rather than information being processed in a domain-specific manner. One domain-general framework is the embedded processes model (Cowan, 1999; Cowan et al., 2020). Unlike the multi-component model, the embedded processes model suggests that WM is a subset of LTM, referred to as activated LTM (Cowan, 1999, 2016; Cowan et al., 2020). Additionally, the embedded processes model proposes that attention can be internally or externally focused towards goal-relevant information in a capacity-limited focus of attention (Cowan, 1993, 1998, 1999). Oberauer's (2002, 2009, 2019) concentric model similarly emphasises a domain-general perspective, with the main structural difference between the embedded-processes and concentric models being that the focus of attention holds only one rather than up to four items.

As alluded in the previous section, there is a great deal of evidence for the domain-general perspective that contradicts the domain-specific conceptualisation of WM. For example, as discussed above, advocates of the domain-specific account of WM have suggested that dual-task paradigms evidence that visual and verbal information is stored in separate domain-specific stores (Cocchini et al., 2002). However, advocates of the domain-general view have argued that dual-task paradigms do not indicate that visual and verbal information are stored in domain-specific stores. For example, in a dual-task paradigm in which participants were required to retain either one set of visual or verbal information, or two sets, Cowan and Morey (2007) observed that recall performance was affected when retaining two sets of information that were of the same domain-type (i.e., visual or verbal). Furthermore, the researchers observed that recall performance was affected when retaining two sets of both visual and verbal information. The researchers argued that these findings may be due to storing more than one set of information, irrespective of the domain's nature. This therefore suggests that a limited attentional resource may be required to keep visual and verbal information active in WM and attention is not domain specific. Arguably, this is similar to the multi-component model's episodic buffer. However, unlike the episodic buffer, the domain-general account suggests that external information is temporarily held active in a subset of LTM without necessary storage in domain-specific stores. This therefore suggests that in a domain-general framework, external information is stored in LTM regardless of its domain and attentional resources, such as the focus of attention, selectively coding goal-relevant information for future processing.

Further evidence in favour of a domain-general conceptualisation of WM has shown that against Logie's (1995) suggestion that WM stores visual information in domain-specific (i.e., within the visual cache and spatial-loop), visual information is stored concurrently (Vergauwe et al., 2009). Across two Experiments, Vergauwe and colleagues (2009) observed that increasing the number of to-be-retained visual and spatial information reduced recall performance, irrespective of whether the increase of to-be-retained items was in the visual or spatial domain. The researchers argued that increasing the number of to-be-retained items reduced the maintenance of in-memory items, and thus attention was divided between encoding and maintenance. Thus suggesting that attentional resources are required to maintain information in WM irrespective of its domain. Advocates of the multi-component model would argue that the divide in attention may demonstrate the distinct central executive and episodic buffer systems, wherein during encoding, the central executive acts as a general processor (through general attention) and during maintenance, the episodic buffer selectively attends to goal-relevant information. However unlike the multi-component model, advocates of the domain-general frameworks would argue that general attention (i.e., focus of attention) would hold information active in activated LTM and selectively attend to goal-relevant information without the requirement of separate attentionally-based systems. Thus it could be argued that parallels can be drawn between the domain-specific and domain-general frameworks of WM, but the storage of information differs between the frameworks.

Moreover, this finding disagrees with Logie's (1995) suggestion that visual information is stored in a passive manner, but spatial information is reliant on a rehearsal mechanism. Furthermore, these findings are aligned with the multi-component model's conceptualisation of the visuospatial sketechpad, in that visual and spatial information are cohesively stored and maintained within it, and the episodic buffer selectively keeps information active. Taking this evidence together, it can be argued that WM may be reliant on additional resources such as attention to maintain information and selectively prioritise goal-relevant information. I will discuss this further in the following sections.

Visual working memory

Visual WM pertains to temporarily stored visual information for goal-orientated processing (Luck & Vogel, 2013). As explained previously, Baddeley and Hitch's (1974) multi-component model was one of the first to provide a framework for understanding visual information processing. In the previous sections, I discussed whether WM may be thought of as domain-specific or domain-general, however in this section I will focus on some of the broader problems within the WM literature concerned with the processing of visual information.

Many of the research questions stem from broader problems in the WM literature, such as how representations are maintained in WM. Seminal research focusing on the field of visual WM has argued that visual information is processed as whole item representations (Luck & Vogel, 1997; Vogel et al., 2001) rather than individual features of an object as had been previously theorised (Treisman, 1999; Treisman & Gelade, 1980; Wheeler & Treisman, 2002). This work from the "slots model" proposed that there is an upper-fixed limit of about three or four visual item representations that can be held "online", in line with Cowan's (2001) hypothesis that WM is limited by the number of items that can be held active within the focus of attention. For example, Vogel and colleagues (2001) used a sequential comparison paradigm, presenting multiple-coloured objects that the participants were required to remember when cued with an object feature. The trials presented between 1-12 coloured objects to understand whether a fixed-item limit would be reached. Experiments 1-10 showed that the participants' memory performance decreased after four or more objects were presented. Vogel and colleagues posited that this is due to individuals storing whole representations of an object rather than aspects of an object, thus implying that visual WM is limited by a fixed item capacity of about four items. Experiments 11-16 supported their observations, such that task performance decreased when items with multiple features (i.e., an item's colour, shape, and spatial location) were presented. The researchers posited that the object representations in visual WM adhere to the fixed item capacity limit. This further implies that the features of an item are automatically encoded and bound in visual WM. In the next section, I will explain the research concerning how item features are bound together.

Following on from this work supporting the "discrete slots" view of visual WM storage, Zhang and Luck (2008) proposed that representations of items can be stored in a fixed number of discrete slots and a resource moderates the processing of the representations within the slots (i.e., slots + resource). Moreover, the researchers proposed that feature representations of a single item may be stored across the discrete slots to allow for a high-quality representation of an item and high precision during item recall (i.e., slots + averaging). In a series of experiments presenting various coloured squares, participants were required to recall the colour of a probed location using a colour response wheel. The researchers observed that when the number of to-be-retained items (i.e., set size) was below 3, the probability that the target was in memory (P_m) was high, but increasing the set size to 6 items reduced P_m. Zhang and Luck (2008) argued that individuals were able to retain high-quality feature representations across the discrete slots using the slots + averaging approach when the set size was low. Conversely, low P_m with the increased set size suggested that the fixed-item capacity of the discrete slots had been reached, and only a limited number of whole item representations were stored, using the slots + resource approach. This may suggest that additional resources are unnecessary until WM capacity is reached, the implications of which I will discuss in the next sections.

Feature binding

Feature binding refers to the integration of aspects of a piece of information, such as the colour, shape, orientation, and location of an object (see Schneegans & Bays, 2019 for a review).

A central question in the feature binding literature concerns whether feature binding is relatively automatic or requires attentional resources.

Some work suggests that additional resources are not required to process feature bindings (Allen et al., 2006; Luck & Vogel, 1997; Vogel et al., 2001). For example, Allen and colleagues (2006) investigated whether individuals could discriminate between different types of objects' features once they are no longer perceptually available in visual WM. The participants undertook four conditions: shape-only, colour-only, shape or colour presentation, and finally coloured shapes. In each condition, four items were presented, and participants were asked to indicate whether an object or colour had changed during the retrieval phase. In Experiment 1, the researchers found that individuals retrieved colour information more accurately than shape information. In addition, the combination of colour and shape condition revealed that there were little statistical differences between the two presented features. The researchers posited that the lack of differentiation in retrieval of the colour and shape objects could suggest that feature binding is an automatic process which does not require great effort. When participants were asked to count backwards during the visual WM conditions in Experiment 2, the results showed that counting backwards during the trials negative affected item retrieval, however, again there was no significant difference between the retrieval of item features. Experiments 3 and 4 replicated the findings of Experiment 2, suggesting that feature binding may be an automatic process. In Experiment 5, the researchers presented the visuospatial features sequentially rather than simultaneously, which hindered item retrieval, particularly in the combination condition. The researchers postulated that presenting additional to-be-remembered items may weaken the already "online" representations. The findings of Allen and colleagues (2006) suggested that visual information may be processed as bound object representations rather than as independent features, thus agreeing with prior seminal visual WM research (Luck & Vogel, 1997; Vogel et al., 2001). Moreover, contrary to Zhang and Luck's (2008) hypothesis of discrete slots + resource, the findings of Allen and colleagues (2006) suggest that the allocation of additional resources is not required for the processing, storage, and maintenance of visual information.

Other research suggests that successful feature binding requires attentional resources. For example, fundamental research in the perception of features of objects proposed that features of an item are processed in their respective feature dimensions, for example, colour features are processed in a colour-store, and shapes are processed in a shape-store (Treisman & Gelade, 1980). While Treisman and Gelade's (1980) framework of visual information processing is thought to be an explanation of the immediate processing of visual information, Treisman and Gelade's framework has largely influenced our current understanding and interpretation of the goal-orientated processing of visual information (see Wolfe, 2020 for a review). Therefore, Treisman and Gelade's (1980) framework must be taken into consideration while examining the processing of visual information in WM and the role of attention. Earlier work by Wheeler and Treisman (2002) suggested that features of an item are encoded, stored, and maintained as individual representations in visual WM, allowing an individual to access the features during active goal-orientated processing. Through a series of experiments, the researchers found that individual features presented did not adhere to a fixed-item WM capacity limit but rather the features were stored in their own specific dimensions. The researchers proposed that a visual WM capacity limit, as proposed in prior research (Vogel et al., 2001), may reflect a single feature dimension's capacity (Wheeler & Treisman, 2002, Experiments 1 and 2) rather than a visual WM capacity limit per se. In line with this feature-based account of information processing, recent hypotheses surrounding the processing of visual information suggest that the quality of maintained single features may depend on the allocation of a shared resource (Bays et al., 2009, 2011; Bays & Husain, 2008). According to this resource view of WM capacity limits, the allocation of a shared resource allows for a visual WM to maintain multiple features by distributing the shared resource towards to-be-prioritised information.

More recently, Li and colleagues (2022) demonstrated that multiple single features and feature bindings may be processed in WM. Specifically, the researchers found that both units of visual WM are stored, but single feature representations are maintained at a higher resolution than feature bindings. This agrees with prior work that suggests that the processing of visual information may depend on the allocation of a shared resource (Bays et al., 2009, 2011; Bays & Husain, 2008), such that storing single feature representations at a higher resolution than feature bindings may rely on attentional resources to keep relevant feature information active in WM. This suggests that attentional resources may be critical to keeping feature representations active in WM. Therefore, it is critical to develop experimental paradigms and powerful analytic tools to investigate whether visual information is processed in an object-based or feature-based manner and what are the underlying components that allow for visual information to be retained in WM.

Paradigms in visual WM

Early feature binding research primarily used change detection paradigms wherein participants try to determine whether the tested array or probe is the same or different to what was presented during encoding (e.g. Vogel et al., 2001). However, it can be argued that change detection paradigms may offer a limited understanding of recall performance because they do not indicate the underlying processes involved with the processing of visual information. As discussed previously, the slots account suggests that visual information is stored in capacity limited slots. In a change detection paradigm, Donkin and colleagues (2014) investigated whether, given the limited capacity of slots (Zhang & Luck, 2008), participants are able to detect whether a change in colour-location binding has occurred. Donkin and colleagues (2014) manipulated the number of colour-location bindings that were required to-be-retained, and during test, participants responded on whether they detected a change. The researchers measured recall through hits (i.e., correct change detection), and false alarms (i.e., a change has occurred, but it was not detected). Donkin and colleagues observed that false alarms increased with the increase in number of items to remember. While this suggests that visual WM may be limited by the number of items it can hold active, it also suggests that change detection paradigms may be less sensitive in determining the cause for false alarms to occur. This suggests that it may be beneficial to understand the underlying components of memory retrieval through alternative paradigms that are able to allow researchers to investigate specific parameters (e.g., guessing).

An alternative method of examining the processes involved in feature bindings is through continuous-report paradigms. During a continuous report paradigm (e.g., Bays et al., 2009; Zhang & Luck, 2008), a participant is presented with an array of objects (usually coloured objects or orientated bars presented around an invisible circle), and following a retention interval, the participant is presented with a response wheel and a probe. The response wheel may be the feature dimension that is being tested, such as a colour wheel and a location-probe to test colour-location bindings (a probe in the location of the to-be-tested item; Bays et al., 2009; Wilken & Ma, 2004; Zhang & Luck, 2008). Participants would then use their mouse to move along the colour wheel and select the most accurate representation of the colour in the probed location. Continuous report paradigms are arguably more advantageous than change detection paradigms because it is possible to extract parameters assumed to underlie retrieval responses. For example, fitting a three-parameter mixture model to recall error allows estimates of recall precision, the probability that the target item was in memory, the probability that the response made was due to a binding error, and the probability that the response made was a guess (Bays et al., 2009; Oberauer et al., 2017). This provides for a critical tool in determining how feature bindings are maintained by providing for a greater understanding than simply the response accuracy from a change detection task (see Appendix A for a further discussion of mixture models).

In particular, the probability of whether the response is a binding error is of interest to understanding the maintenance of feature bindings in visual working memory. Binding errors or ‘swap errors’ refer to mistakenly recalling the feature of a presented but non-tested item (Bays et al., 2009; Schneegans & Bays, 2019). For example, Emrich and Ferber (2012) found that binding errors increase if the spatial-locations of the to-be-remembered colours are close together. Further work has supported this finding, such that individuals are more likely to make binding errors if there is a similarity between the spatial-locations of the to-be-remembered items (Bays, 2016). Moreover, in a continuous-report paradigm, Rajsic and Wilson (2014, Experiment 1) demonstrated that binding errors occur more often while recalling a location feature compared to a colour feature. Rajsic and Wilson (2014) argued that location and nonlocation features (i.e., colour) are coded in conjunction in visual maps, such that the code of a colour feature may be mapped on to its respective spatial location. Furthermore, the researchers argued that tuning a visual map while recalling a location feature may result in tuning multiple visual maps at the same time, thus binding errors occur between to-be-recalled locations. On the other hand, the researchers argued that fewer codes are available for nonlocation feature storage in visual maps; therefore, fewer binding errors may occur during the retrieval of a colour feature. This hints that there may be differences between the processing of location and non-location features, in line with prior research in perception (Treisman, 1977; Treisman & Gelade, 1980; Treisman & Zhang, 2006; Wheeler & Treisman, 2002), and with more recent research in visual WM that argued that location features may be prioritised over non-location features (Schneegans & Bays, 2017). Thus, it can be argued that using continuous-report paradigms may be a more ideal way to investigate feature binding than change detection paradigms.

In summary, visual WM feature binding is heavily researched and there is an ongoing debate as to how visual information is processed in WM. There appears to be a divide in the current literature. One view argues that visual information is processed in a feature-based manner (Rajsic & Wilson, 2014; Schneegans & Bays, 2017; also see Swan & Wyble, 2014; Treisman, 1977, 1999; Treisman & Gelade, 1980; Wheeler & Treisman, 2002). An alternative view suggests that visual information is processed in an object-based manner (Allen et al., 2006, 2012; Luck & Vogel, 1997; Vogel et al., 2001). While earlier research has primarily examined visual WM through change detection paradigms (e.g., Luck & Vogel, 1997; Vogel et al., 2001), more recent research has used continuous-report paradigms to allow for mixture models to be applied to the recall responses (e.g., Adam et al., 2017; Bays et al., 2009; Zhang & Luck, 2008). In the next sections, I consider the research concerning the prioritisation of goal-relevant information in visual WM.

Does attention aid the processing and prioritisation of visual information?

A great deal of research suggests that attention is important to the encoding and maintenance of information in WM (Astle et al., 2012; Camos & Barrouillet, 2014; Griffin & Nobre, 2003; Gunseli et al., 2015; Hajonides et al., 2020; Li & Saiki, 2015). From a multi-component perspective, Hitch and colleagues (2020) proposed that the processing of visual features initially relies on the allocation of external attention (i.e., attention can be directed towards incoming information), followed by the formation of object files in the visuo-spatial sketchpad. The researchers suggested that the object files are then processed by the episodic buffer by means of internal attention (i.e., attention that is directed towards in-memory representations). Hitch and colleagues (2020) argued that the processing of object files (and their respective features) is governed by the central executive allocating attentional resources towards to-be-refreshed information. This allows for recent visual information to be kept active.

This acknowledgement of the role of attention in the multi-component model of WM is consistent with WM research from the last five decades that has argued that attention is an important factor in keeping information active in memory (see Oberauer, 2019 for a review). As previously discussed, the domain-general account of WM (such as the embedded-processes and concentric models; Cowan, 1999; Cowan et al., 2021; Oberauer 2001) argues that general attention (i.e., the focus of attention) is critical for keeping information active. Thus it may be argued that the processing of visual information may be reliant on attentional resources. Seminal research in the perception of visual features posited that the binding of visual features requires attentional resources (Treisman, 1977, 1999; Treisman & Gelade, 1980). Moreover, it has been posited that the attentional resources required for the binding of visual features may selectively prioritise location-based features over non-location features (Treisman & Zhang, 2006; Wheeler & Treisman, 2002). Further fundamental research investigating attention towards visual features in WM has used location-based prospective cues (pre-cues; i.e., the target is presented before the array of to-be-remembered items) and retrospective cues (retro-cues; i.e., the target feature is presented after the array of to-be-remembered items, and after a brief interval) that guide attention towards to-be-prioritized information (Griffin & Nobre, 2003). Griffin and Nobre (2003) demonstrated that guiding attention via pre-cues and retro-cues improved recall accuracy compared to neutral (no-cue, baseline) conditions. Moreover, the researchers' results suggested that the valid cue-types improved recall accuracy compared to invalid cue-types. Griffin and Nobre's (2003) results thus suggest that guiding attention via a location-based pre-cue and retro-cue, in comparison to a neutral cue, allows for spatial-based attentional resources to selectively prioritise goal-relevant information. Arguably this suggests that location-based features may be critical in allowing attention to be orientated towards goal-relevant information regardless of whether the feature is pre- or retro-cued. Retro-cues however have since become a productive way of understanding the role of attention in visual WM (see Souza & Oberauer, 2016, for a review).

Importantly, Griffin and Nobre's findings have raised multiple research questions, such as whether the nature of the retro-cue important in its effectiveness to prioritise feature information. Recent research has demonstrated that the nature of the retro-cue does not impede on recall performance. In a series of experiments, Arnicane and Souza (2021) demonstrated that a retro-cue benefit occurs in target memory (i.e., the likelihood that the to-be-retained feature is in-memory during recall), regardless of the nature of the retro-cue (e.g., orientation-cue, colour-cue, or shape-cue). Moreover, in a series of experiments manipulating the presentation of pre- and retro-cues, Li and Saiki (2015) found that colour pre-cues were more effective in modulating attention than colour retro-cues, relative to a no-cue baseline condition. However, location retro-cues were more effective in guiding attention than location pre-cues, relative to a no-cue baseline condition. Taken together with the findings of Arnicane and Souza (2021), it can be argued that guiding attention via pre-cues and retro-cues strengthen memory traces for recall, with some potential differences in the efficacy between pre-cues and retro-cues. In the next section, I will address further research regarding whether outdated (irrelevant) information is still accessible in visual WM while prioritising goal-relevant information.

What happens to outdated (irrelevant) information during the prioritisation of goal-relevant information?

Although it is critical to understand the underlying mechanisms involved with prioritising goal-relevant information in visual WM (by means of attention), it is also critical to understand what happens to outdated (irrelevant) information while we prioritise goal-relevant information. One way to examine if outdated (irrelevant) information is kept active in WM is by investigating a core function of WM known as WM updating. WM updating refers to the ability to manipulate and alter goal-relevant information (Lewis-Peacock et al., 2018). Arguably, WM updating is a core function of WM as we need to manipulate and alter processed information to achieve our goals. For example, while shopping for a new t-shirt, one may store features of a t-shirt such as its colour and pattern. One may continue browsing for other t-shirts and find a similar item with the same pattern as the in-memory t-shirt, but the new item has a different colour. One actively prioritises the new (updated) information and the outdated colour information is now irrelevant. This example demonstrates that during WM updating, new features are prioritised and outdated features are no longer required. However, is outdated information still accessible in WM, and does outdated information impact the prioritisation of relevant information? Thus, it is of interest to examine whether outdated feature information is still accessible in WM while one prioritises new features, or whether the outdated features are removed and now longer accessible.

Initial WM updating research posited that irrelevant information may be suppressed while new goal-relevant information is prioritised (Palladino et al., 2001). Across five experiments, Palladino and colleagues (2001) investigated whether there are differences between one's ability to comprehend language (poor or good comprehends), and one's ability to prioritise updated information. In their set of experiments, the researchers orally presented various lists of words and the participants were required to prioritise the task-relevant information. In Experiments 2 and 3, the researchers manipulated whether participants were presented with concrete (animals) and abstract nouns. The participants were instructed to prioritise the concrete words and to ignore the abstract words; this was followed by a memory recall task, wherein the participants recalled *all* the presented items. Across their set of experiments, Palladino and colleagues found that the ability to update WM depended on the participants' language comprehension. Furthermore, the researchers posited that prioritising goal-relevant verbal information required the suppression of irrelevant information; this was particularly evident in Experiments 2 and 3. Arguably, this suggests that while prioritising new goal-relevant information, outdated information is still accessible, but its memory traces are not as potent as the prioritised new information.

Oberauer (2002) posited that outdated (irrelevant) information may still be accessible in WM, despite the focus of attention prioritising task-relevant information. This work suggests that prioritising new information during WM updating may rely on the removal of outdated information from WM, rather than its suppression. In a recent review, Lewis-Peacock and colleagues (2018) proposed that the removal of outdated information may occur temporarily (i.e., outdated information is no longer required in WM, but it is accessible for future processing) or permanently (i.e., outdated information is removed from WM, and it is inaccessible). The researchers argued that the removal of outdated information occurs through the removal of an item's content, context, or content-context binding. Earlier work demonstrated that during WM updating, outdated information is removed by the transfer of old item-position bindings to allow for new item-position bindings to occur (Ecker, Oberauer, et al., 2014). Across three experiments, Ecker and colleagues (2014) presented participants with three rectangular frames in a row that remained displayed throughout the experiments. In each frame, a consonant letter was presented for 2000 ms. Following this, participants were briefly presented the blank frames, and after a short interval, one of the frames was highlighted to indicate that the consonant in that frame is to-be-updated. The to-be-updated consonant was then presented in the highlighted frame for 5000 ms. The participants were required to update the letter in their WM and to remove the outdated information. Participants were instructed to press the spacebar on the keyboard once the information has been updated. The presentation of the blank frames followed by the highlighted frame and the updated letter varied between 1 and 21 times. During the retrieval phase, the participants were required to recall and type each new (updated) letter in the randomly probed positions of the rectangular frames.

In Experiment 1, the researchers manipulated the presentation time of the blank, highlighted to-be-updated frames and the presentation time of the new (updated) letters. In Experiment 2, the researchers presented numbers rather than letters and manipulated whether the numbers were proximally close or distant (e.g., proximally close: 11 & 13, proximally distant: 11 & 45). In Experiment 3, the researchers examined whether WM capacity predicted WM updating capabilities by including a battery of WM capacity tasks in their experiment. Across the three experiments, the researchers observed that actively updating WM required removal of the item-position bindings of outdated information by prioritising new (updated) information at high speeds. Arguably, contrary to Palladino and colleagues' (2001) suggestion that outdated information is suppressed in WM, Ecker and colleagues' (2014) work suggests that actively removing outdated information in a bound representation is not effortful and only new bindings are prioritised in WM.

Similarly to Ecker and colleagues' (2014) suggestion that outdated information is removed by the removal of item-position bindings, an earlier WM updating account posited that prioritising new information and removing outdated information can occur locally (only the outdated information is removed from a binding, and the remaining contents of WM are not updated; i.e., outdated features within a binding are removed and replaced by updated features), and globally (the entire contents of WM are updated, i.e., all features and bindings are removed and replaced with updated features and bindings; Kessler & Meiran, 2008). Kessler and Meiran (2008) proposed that WM updating initially occurs through local updating mechanisms, but in turn global updating occurs to form a new cohesive representation of to-be-retained information. This suggests that local and global WM updating may reflect the core process of prioritising task-relevant information. An investigation into whether visual WM is locally or globally updated showed that partially-updating colour-location bindings (i.e., updating some, but not all of the visually presented items) requires local WM updating (Kessler et al., 2015). By manipulating whether a repeated item-location (the location of a colour that has not updated) or an updated item-location (the location of a colour that has updated) was probed, the researchers found that recall accuracy was best when probed with a repeated-item location relative to an updated-item location. This suggests that during local updating, outdated information may still be accessible, and it has not been permanently removed. However, it is unclear whether the availability of outdated information in WM limits the prioritisation of new information. Therefore, it can be argued that a closer examination of the prioritisation of new information and the removal of outdated information is required. Moreover, it is of interest to examine WM updating of visual information as this may provide a clearer indication of the removal of outdated information, due to the dissimilarities between feature dimensions (i.e., colour, shape, or location feature-types).

Summary and Overview of Experiments

There is a great debate surrounding the processing of visual information in our immediate memory. A domain-specific perspective suggests that visual information may be processed in its domain-specific store (e.g., visual-spatial sketchpad; Baddeley and Hitch, 1974; also see Baddeley, 2012). In contrast, a domain-general perspective suggests that information of any domain is processed by focusing one’s attentional resources towards to-be-remembered information, and thus initiating activated LTM (Cowan, 1993, 1998, 1999; Cowan et al., 2020; Oberauer, 2002). Furthermore, there is still some uncertainty as to whether visual information is processed as whole-item representations, or whether visual information may be processed as single features and feature bindings (see Schneegans & Bays, 2019). This raises the question: How is goal-relevant feature information processed and prioritised in visual WM? One aim of this thesis is to examine the role of attention during the prioritisation of visual information in WM. Secondly, this thesis examines whether non-prioritised (outdated, or irrelevant) visual information is accessible during WM updating. It has been argued during WM updating, new (updated) information is prioritised but outdated information is no longer required (Ecker, Oberauer, et al., 2014; Lewis-Peacock et al., 2018). To do this, this thesis aims to address how single features and feature bindings are processed in WM to allow for the prioritisation of goal-relevant information. I briefly summarise the findings of my chapters below.

In Chapter 2, across two experiments, it was investigated what the role of attention is during feature binding. In Experiment 1, two types of attention were disrupted: peripheral- and central-attention. In Experiment 2, attention was guided via retrospective (retro-cues). In Experiment 1, it was observed that disrupting attention did not increase binding errors, irrespective of the disruption-type or the target feature. In Experiment 2, it was observed that guiding attention did not reduce binding errors, irrespective of the nature of the retro-cue or the target type. However, it was observed that disrupting attention decreased target memory if the target feature was location-based. It was also observed that guiding attention improved target memory irrespective of the nature of the retro-cue or the nature of the target feature.

Chapter 3 followed-up on Experiment 1 of Chapter 2 by reducing the number of to-be-remembered items from six to four. It was observed that disrupting peripheral- and central-attention did not increase binding errors, relative to a no-disruption baseline. Against the findings of Experiment 1 in Chapter 2, disrupting peripheral- and central-attention did not reduce target memory whilst the number of to-be-remembered items were reduced. However, further analyses of the recall errors of Chapter 3 in comparison to the recall errors of Experiment 1 of Chapter 2 indicated that task performance improved while the number of to-be-retained items was reduced.

In Chapter 4, it was investigated whether location-based features are privileged in visual working memory, in comparison to non-location features (such as colour). An adapted paradigm from Experiment 2 in Chapter 2 was used to answer our research question. Chapter 4 compared the baseline conditions (uncued retention interval) and the retro-cue conditions relative to the baseline condition, respectively. It was observed that there were no credible effects between the probe-target types (colour-location, location-colour) in the baseline conditions. Guiding attention via colour and location retro-cues did not reduce binding errors, relative to the baseline condition. However guiding attention via colour and location retro-cues increased target memory irrespective of the nature of the target, and in addition, guessing was reduced. Thus suggesting that there are no differences between the processing and attentional prioritisation of location and non-location feature dimensions.

Chapter 5 investigated whether local working memory updating relies on temporary removal of irrelevant information (i.e., irrelevant information is momentarily withdrawn from working memory, but it is still accessible) or whether irrelevant information is permanently removed (i.e., it is no longer active in working memory and no longer accessible). Chapter 5 adapted a prior change detection paradigm, to measure recall on a continuous scale, using a delayed estimation paradigm. Two hierarchical Bayesian three-parameter mixture models simulating temporary (Model 1) and permanent removal (Model 2) were applied to the data. Model 1 indicated that binding errors increase while more than one item is to-be-updated, relative to the baseline condition. But in Model 2, it was observed that binding errors were similar irrespective of the number of items that were required to be updated. Further analyses indicated that the model simulating temporary removal best fit the data.

Chapter 6 followed-up on Chapter 5 by changing the feature type that was to-be-updated from a colour feature to a shape feature. Chapter 6 applied the same statistical analyses to the responses. In both models fitted, it was observed that binding errors were similar irrespective of the update condition, or the probe location. Comparing model fit, it was observed that the model simulating temporary removal best fit the data. Moreover, by comparing the recall errors of Chapter 5, relative to the recall errors of Chapter 6, it was observed that recall performance declined while shape features were required to-be-updated in comparison to updating colour features. Thus suggesting that it is more difficult to locally update shape features compared to colour features in working memory.

Chapter 2

Examining the role of attention during feature binding in visuospatial working memory

This chapter is a registered report that was published with Nicholas R. Cooper and Vanessa M. Loaiza at Attention*, Perception & Psychophysics*: Goldenhaus-Manning, D., T., Cooper, N., R., & Loaiza, V., M. (2023). Examining the role of attention during feature binding in visuospatial working memory. Attention, Perception, & Psychophysics, 1–12. https://doi.org/10.3758/s13414-023-02655-y. I contributed to this article by generating the project research question and hypotheses, writing the initial drafts of the registered report for submission (stage 1 review), formatting the stimuli used within the task, aiding with the programming of the experiments, data collection, aiding with data analysis, and disseminating the findings in the form of a registered report which was submitted for stage 2 review.

Fundamental questions of human memory concern how individuals are able to focus their attention and hold in mind no-longer perceptually available items and how features of that item (e.g., colour, shape, orientation and relationship to other items in mind) are integrated together into a cohesive representation. Working memory (WM) has been an instrumental concept for understanding these questions given its ascribed role in maintaining and manipulating information in service to goal-related cognition. For example, in a typical visuospatial WM study, participants are often asked to briefly maintain an array of multi-feature objects (e.g., red circle, green triangle, blue square) to be immediately retrieved thereafter (e.g., via change detection, determining whether the probe of a red triangle was presented). A prominent yet still unresolved question in the literature concerns whether integrating the features of these objects, or feature binding, relies on attention or is relatively automatic (Allen et al., 2006; Baddeley et al., 2011; Elsley & Parmentier, 2009; Hitch et al., 2020; Vogel et al., 2001; Wheeler & Treisman, 2002). In the current study, we investigated this issue through two predominant methods of manipulating attention in the WM literature: disrupting attention and guiding attention.

Disrupting Attention and Feature Binding

A great deal of prior research has followed the logic that if feature binding requires attention, then an attention-demanding task should more strongly disrupt the maintenance and recall of bindings (e.g., remembering a red circle) compared to individual features (e.g., remembering red and circle individually). Much of this research has demonstrated that attentionally-demanding tasks (e.g., judging tones, counting backwards) similarly impair the recall of individual features and their bindings, thus suggesting that feature binding may not require additional attentional resources (Allen et al., 2006, 2012; Langerock et al., 2014; Morey & Bieler, 2013; Vergauwe et al., 2014). However, other research has reported contradictory findings. For example, some work has suggested that attention-demanding tasks do indeed impair feature binding disproportionately compared to individual features, thus suggesting that feature binding in WM requires additional attentional resources (Brown & Brockmole, 2010; Elsley & Parmentier, 2009; Zokaei et al., 2014). What may cause such a discrepancy?

First, the nature of the to-be-bound features may be an important factor in whether feature binding requires attention. A great deal of the prior work supporting the notion that feature binding is automatic has used colour-shape bindings. According to Ecker and colleagues (2013), such *intrinsic* feature bindings (i.e., information belonging to an object) can be distinguished from *extrinsic* feature bindings (i.e., relational, contextual information of an object, such as its spatial location). This distinction between intrinsic and extrinsic feature binding has also been referred to as conjunctive versus relational binding, respectively (Kirmsse et al., 2018; van Geldorp et al., 2015). Ecker and colleagues suggested that binding of intrinsic features (e.g., colour, shape, orientation) may be automatic in WM, whereas binding of extrinsic features (e.g., spatial/temporal context) may require additional, effortful attentional resources. Consistent with this assertion, some prior work has indeed shown a role of attention for extrinsic bindings (Ecker et al., 2013; Elsley & Parmentier, 2009), whereas attentionally-demanding tasks have no differential impact on intrinsic feature binding (e.g., colour-shape bindings) compared to retaining the individual features, as explained previously (e.g., Allen et al., 2006, 2012). Conversely, other work has suggested that visual interference can impact visual WM (Teng & Kravitz, 2019), especially intrinsic feature bindings (Ueno et al., 2011). Thus, it may be the case that disrupting central attention specifically impairs extrinsic binding but not intrinsic binding, whereas disrupting peripheral, visual attention impairs intrinsic binding for visuospatial features. This distinction may occur if extrinsic features are prioritized and provide the foundation for intrinsic information to be automatically encoded. For example, Schneegans and Bays (2017) demonstrated that intrinsic features may rely on spatial-temporal context in order to become feature-bound representations in WM (see also Pertzov & Husain, 2014).

A second potential source of discrepancy concerns the measurement of feature binding. Much of the cited previous research has used change detection, wherein participants are presented with a test probe that may be the same as what had been presented or a lure, such as a recombination of features. A potential problem with this approach is that using observed performance may only offer a coarse measure of feature binding that is conflated with extraneous processes like guessing.

For example, a participant may correctly identify a red circle had been presented, but it is unclear whether this indicates correct intrinsic binding or other processes, such as guessing. Instead, measuring recall using a continuous scale (e.g., recalling the precise colour or location of a probe) allows decomposition of the participants' observed recall error (e.g., the deviation of the recalled colour from the true target colour) according to a mixture of underlying components. Most relevantly, Bays and colleagues' (2009) model assumes three sources of error: memory for the target item with a certain precision, memory for one of the presented but unprobed items (henceforth, binding errors), or random guessing when no memory is available. Previous work applying this mixture model to recall has indicated that binding errors may be more common when recalling extrinsic versus intrinsic features (Schneegans & Bays, 2017) or when attentional demands are high (Zokaei et al., 2014). For example, Zokaei and colleagues (2014) showed that varying the demand of an unrelated visual search task during the retention interval of a visuospatial WM task primarily impacted binding errors, whereas precision and random guessing were largely unaffected. Thus, the parameter estimate representing binding errors when applying a mixture model may be more sensitive than observed performance in determining the role of attention in feature binding.

Guiding Attention and Feature Binding

Alongside the substantial number of studies that have explored the role of attention in feature binding via disrupting attention, another approach is to manipulate attention through instructions to prioritize some memoranda. For example, presenting retro-cues during the retention interval to indicate the to-be-tested item of the previously presented stimulus array provides a means of understanding how no-longer perceptually presented information is brought into the focus of attention (see Souza & Oberauer, 2016 for review). Retro-cue studies often show a retro-cue benefit, such that retro-cues improve visuospatial WM performance compared to a no-cue or neutral-cue baseline, thereby demonstrating the benefits of attention in WM. Furthermore, the retro-cue benefit seems specific to reducing binding errors (Souza, 2016), thus evidencing the importance of attention to facilitate binding in WM. There is some evidence that the retro-cue benefit is weaker for intrinsic (e.g., shape) compared to extrinsic (e.g., spatial-location) retro-cues (Arnicane & Souza, 2021; but see Heuer & Schubö, 2016). This may indicate that intrinsic retro-cues are less effective overall than extrinsic retro-cues or perhaps are just less effective for extrinsic bindings compared to intrinsic bindings, but this has not yet been explicitly tested. Coupled with the aforementioned research pointing to a potential distinction between the effects of disrupting central versus peripheral attention on extrinsic and intrinsic binding, respectively, these findings suggest a dissociation in how guiding attention may impact different types of feature binding in WM. Explicitly investigating this dissociation would provide novel insight into the long-standing theoretical puzzle regarding how representations are established in WM.

Current Study

The current study clarifies the role of attention in feature binding in WM by investigating whether manipulating attention via disruption (Experiment 1) and retro-cues (Experiment 2) differently impacts the maintenance of intrinsic and extrinsic feature bindings (Table 1). Both experiments required participants to maintain a set of multi-feature objects (i.e., different-coloured shapes) presented in random locations around an invisible circle, followed by probed recall of one of the items. Participants were prompted to recall one of the features based on one of the two remaining features (e.g., a shape probe may prompt recall of colour, with the probe displayed at the centre of the screen [i.e., without spatial information]). Experiment 1 disrupted different types of attention (central versus perceptual) whereas Experiment 2 manipulated focused attention via different types of retro-cues (intrinsic [colour, shape], extrinsic [spatial]). Experiment 1 closely followed the experimental procedure of Souza and Oberauer (2017, Experiment 1A), and Experiment 2 closely followed the experimental procedure of Souza (2016). Both experiments had a no-disruption/no-cue baseline condition. Recall error was fit with a hierarchical Bayesian three-parameter mixture model (Bays et al., 2009; Oberauer et al., 2017) to estimate latent cognitive parameters underlying observed memory performance, most relevantly, the parameter reflecting binding errors.

We hypothesized that if extrinsic and intrinsic feature binding are distinguishable, then a central-attention demanding task should increase binding errors when recalling extrinsic features, but not intrinsic features, whereas a peripheral-attention demanding task should increase binding errors when recalling intrinsic features, but not extrinsic features (Experiment 1). Furthermore, extrinsic retro-cues should reduce binding errors compared to a no-cue baseline, whereas there should be no impact of intrinsic retro-cues on binding errors when recalling extrinsic features (Experiment 2). However, if there is no distinction between extrinsic and intrinsic binding and both are automatic, then binding errors should be similar regardless of the nature of the attentional demand (Experiment 1) or retro-cue (Experiment 2).

Table 1

Hypotheses and Predictions for Each Feature Binding Type.

Feature binding type	Hypotheses	Experiment (E) Predictions
Extrinsic	Central attention may be required to selectively maintain extrinsic bindings	(a) Disrupting central attention should selectively increase extrinsic binding errors but not intrinsic binding errors relative to a no-disruption baseline (E1).
	Benefits of guiding attention may be specific to extrinsic binding.	(b) Guiding attention via extrinsic retro-cues should reduce extrinsic binding errors but not intrinsic binding errors relative to a no-cue baseline (E2).
Intrinsic	Domain-specific, peripheral attention may be required to selectively maintain intrinsic features.	(c) Disrupting peripheral attention should selectively increase intrinsic binding errors but not extrinsic binding errors relative to a no-disruption baseline (E1).
	If intrinsic binding errors are reduced during the presentation of an intrinsic retro-cue, then attention is required for intrinsic feature binding.	(d) Guiding attention via an intrinsic retro-cue should reduce intrinsic binding errors relative to a no-cue baseline (E2).

Experiment 1

Method

Participants. We recruited 24 unique participants per experiment from Prolific in exchange for UK£7.50 per hour of participation. Participants in both experiments were British, aged 18-35, and had normal or corrected-to-normal hearing and vision and normal colour vision. Participants in both experiments were required to pass an initial colour blindness test; those who did not pass it were not allowed to proceed further in the experiment.[1] Participants provided informed consent and were fully debriefed in both experiments. All experiments in this thesis were approved by the

[1] Note that we did not define what we meant by "pass" at stage 1. The Ishihara colour blindness test (e.g., http://www.colourvisiontesting.com/ishihara) that we used typically requires 100% pass rate (e.g., Loaiza & Souza, 2019) and thus we used this criterion for participants to proceed further with the experiment.

University of Essex ethics committee and are in accordance with the Helsinki ethical guidelines (World Medical Association, 2013). Table 2 shows the final sample details, with details of the exclusions explained further on in the next sections. The anonymized raw data of all the participants are available on the Open Science Framework (OSF; https://osf.io/jr3eh/).

Table 2

Sample Details and Exclusions

	Exp. 1	Exp. 2
Total N attempted	61	52
N exclusions:	37	28
1. Failed to pass the first colour blindness/visual/auditory screening phase*	30	10
2. Restarted in the middle of the experiment*	0	4
3. Assigned to a counterbalance order that was already complete*	1	10
4. Incomplete data (e.g., from leaving the experiment)*	6	8
Final N for analysis after exclusions	24	24

* Note that these may not sum to the total N excluded given that participants could fall under more than one category for exclusion.

We determined the sample size and the number of trials per condition by simulating 150 experiments based on the parameter estimates derived from fitting a hierarchical Bayesian three-parameter mixture model to the raw data of Souza and Oberauer (2017, Experiment 1A) and Souza (2016), whose designs we closely follow in Experiments 1 and 2, respectively. For the sake of brevity, the rationale, analysis scripts, and results of the simulations for both Experiments 1 and 2 can be found on the OSF.

Materials and Procedure. The stimuli and the open-source scripts for all the experiments are available on the OSF, and a short example can be tried at this link: https://tinyurl.com/RegReportAPP. Both experiments were conducted online through lab.js (Henninger et al., 2021) hosted on the JATOS server Mindprobe (https://jatos.mindprobe.eu; Lange et al., 2015). After a brief demographic questionnaire, participants took part in two practice phases of the visual and auditory tasks, respectively, that served as the later attention disruption in the critical task. Participants who did not pass either practice phase were not allowed to proceed further as it suggested that they did not have sufficient visual or auditory abilities to complete the experiment.[2] Thereafter, participants completed three blocks of a visuospatial WM task, with four practice trials preceding each block and 150 critical trials presented per block (50 of each probe-target type, randomly intermixed)[3]. Each trial began with a fixation cross presented for 500ms followed by an array of six coloured shapes simultaneously presented in random locations around an invisible circle for 1000ms. The colour, shape, and location features of each item were randomly sampled from 360^{o} of continuous values, with a minimum of 15^{o} of separation in each feature domain from the other memoranda in the array: The colours were sampled along a circle in the CIELAB colour space (with L = 70, a = 20, b = 38, and radius = 60), the shapes were drawn randomly from a shape wheel (Li et al., 2020), and the locations were drawn randomly along an invisible circle (radius = 150[4]).

[2] Note that we did not define what we meant by "pass" at stage 1, and thus we decided that at least 80% correct was required given that this is a typical exclusion criteria in dual-task WM experiments (e.g., Ricker & Vergauwe, 2022).
[3] Note that due to a server upload error, there were several participants for whom some trials were missing at random (*N* = 1 and 3 in Experiments 1 and 2, respectively). We decided to include participants who had at least 80% of trials per cell of the design.
[4] Note that we had written radius = 40 here at stage 1, but this was a typo. A radius of 150 is typical (e.g., Loaiza & Souza, 2019) and allowed better spacing of the memoranda in the array.

Following Souza and Oberauer (2017, Experiment 1), the nature of the retention interval (2500ms in total) varied according to the attention disruption manipulation: During the no-disruption block, the retention interval remained blank with the fixation cross at the center of the screen. During the peripheral-attention disruption block, the fixation cross altered its shade from white to light grey for 100ms during 50% of the trials at least 500ms after the offset of the memory array and 900ms before the onset of the retrieval phase; participants were instructed to detect whether the fixation cross changed its shade by pressing the spacebar (Figure 1). During the central-attention disruption block, participants were asked to indicate whether two successively presented tones (75ms each) were of a lower (600 Hz) or higher (675[5] Hz) pitch by pressing the left- and right-arrow keys, respectively. The first tone was presented 500 ms after the offset of the memory array, with participants allowed 925ms to respond before the second tone was presented with a 925ms response period.

During the retrieval phase, participants were probed with one feature from one item of the memory array to prompt recall of one of the item's two remaining features. For each probe-target type, retrieval occurred along a continuous colour, shape, or location wheel (depending on the to-be-recalled target feature), wherein participants used a mouse to click along the corresponding wheel[6]. For the intrinsic-intrinsic probe-target condition, the participants were probed with either the colour (presented as a circular dot) or shape (presented in dark grey) at the centre of the screen in order to recall the item's corresponding shape or colour, respectively. For the extrinsic-intrinsic

[5] Note that we had written 610 Hz in the revised stage 1. However, all the participants failed the first screening phase, suggesting that it was impossible to distinguish tones presented at 600 and 610 Hz. We conducted an additional pilot experiment (see OSF and Online Supplementary Materials (OSM)) to determine that 600 and 675 Hz, with additional opportunities for feedback and practice during the screening phase, would make the task challenging albeit still possible for participants.

[6] Recall error was slightly underestimated by about 3° in Experiment 2 due to a programming error. Given that this error was unsystematic, the general pattern of results from Experiment 2 is unlikely to be affected.

probe-target condition, participants recalled the colour or shape feature when probed with the location of the corresponding item (i.e., a dark-grey circular disk appearing in the location of the probed item). Finally, for the intrinsic-extrinsic probe-target condition, participants recalled the location of the probed item with either the colour (presented as a circular dot) or shape (presented in dark grey) at the centre of the screen. The recall attempt was unspeeded, and the instructions emphasized that participants should prioritize accuracy over speed in their responses. After the practice trials and every 10 test trials, participants received feedback about their average recall accuracy (expressed as a percentage of their mean reproduction error, i.e., 100 – 100 * mean error/180)[7] and, depending on the block, their average accuracy on the disruption task. There was an inter-trial interval of 1000ms followed by a screen that said "Ready?" to which participants pressed the spacebar to proceed to the next trial. Participants were offered a break twice during each block, after every 50 trials.

[7] For several participants in Experiment 2, the recorded recall error of the shape target was systematically off by 90°. This was corrected during analysis and did not affect the experiment except that these participants' feedback was slightly incorrect.

Figure 1

Example Trial Sequence from Experiment 1 that Varied Attention Disruption (None, Peripheral, Central) and Probe-Target Type (Intrinsic-Intrinsic, Extrinsic-Intrinsic, Intrinsic-Intrinsic).

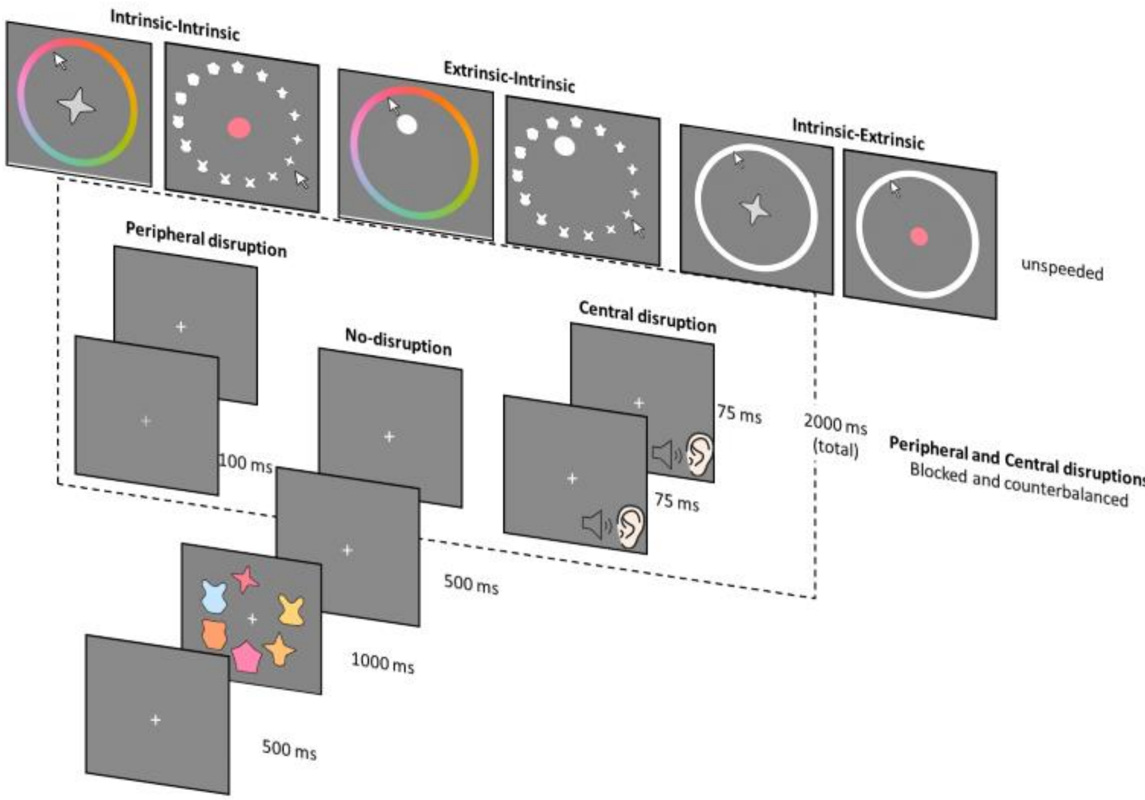

Design. The experiment followed a 3 (attention disruption: none, peripheral, central) x 3 (probe-target type: intrinsic-intrinsic [colour-shape, shape-colour], extrinsic-intrinsic [location-colour, location-shape], intrinsic-extrinsic [colour-location, shape-location]) within-subjects design. The attention disruption manipulation was blocked and counterbalanced [8] across participants, with the nature of the probe-target type varying randomly within each block.

[8] Note that there were instances in both experiments where participants were inadvertently assigned to a counterbalance order that had already been completed (e.g., in Experiment 1, six counterbalance orders required four participants each to equal 24 total participants, but one additional participant was assigned to one of the orders). In these instances, these additional participants were excluded from analysis in both experiments (see Table 2).

Data Analysis. The analysis scripts to reproduce the results for all the experiments are available on the OSF. For each experiment, observed recall error data collected was fit with a three-parameter hierarchical Bayesian mixture model (Loaiza & Souza, 2019; Oberauer et al., 2017). The model assumes that the distribution of observed responses reflects the contributions of (1) the probability that the tested feature is held within WM with a (2) specific precision, and (3) the probability of misbinding or (4) guessing when the participant has not stored the information in WM. Our hypotheses pertained to binding errors, or the probability of recalling non-target but presented features from the array. The model was fit using rjags (Plummer, 2016) via Markov Chain Monte Carlo sampling. We verified good convergence and conducted posterior predictive checks to ensure appropriate model fit. We report posterior estimates of each parameter in Figures 2, 4, and 5, and the posterior differences between the conditions for each parameter (and their 95% HDIs) in Tables 3 and 4.

Results and Discussion

Table 3 and Figure 2 summarize our results. When recalling the extrinsic feature (i.e., the intrinsic-extrinsic condition), binding errors were *lower* in the central (estimated difference = -0.32 [-0.51, -0.13]) and peripheral (estimated difference = -0.22 [-0.43, -0.01]) conditions compared to the no-disruption baseline. Closer inspection of Figure 2C reveals that this result may reflect an artifact of particularly high binding errors at baseline in the intrinsic-extrinsic condition. Furthermore, there was no credible effect of disrupting central attention when probed with the extrinsic feature (i.e., the extrinsic-intrinsic condition; estimated difference = 0.00 [-0.26, 0.28]). These results thus conflict with our first hypothesis, instead showing that disrupting central attention did not have any specific detrimental effect on binding errors of extrinsic features.

Figure 2

Posterior Parameter Estimates of the Bayesian Hierarchical Mixture Model for the Probability of Target Memory (A), Memory Precision (B), Probability of a Binding Error (C) and Probability of Guessing (D) in Experiment 1. Larger Dark Circles Indicate Group Means, Smaller Faded Circles Indicate Individual Means, and the Error Bars Show the 95% HDIs of the Posterior.

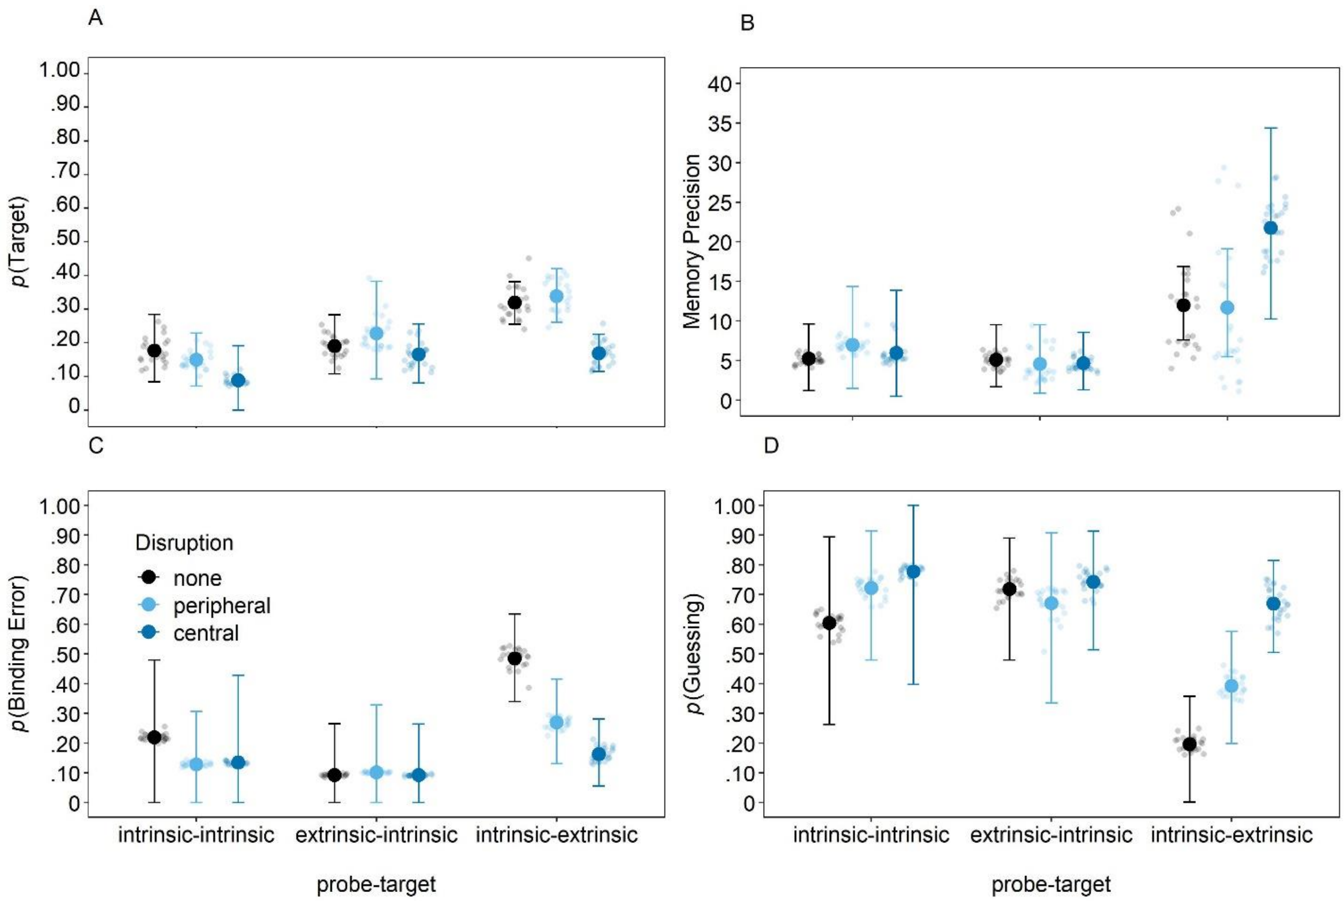

Table 3

Summary of Mean Differences [and 95% HDIs] in Each Memory Parameter Across Experiments.

Exp.	Probe-target type	Baseline/None vs.	*P*(Target)	*P*(Binding error)	*P*(Guessing)	Precision
1	Intrinsic-intrinsic	Central	-0.09 [-0.24, 0.06]	-0.09 [-0.48, 0.32]	0.17 [-0.32, 0.70]	0.73 [-7.46, 10.47]
		Peripheral	-0.03 [-0.16, 0.11]	-0.09 [-0.44, 0.24]	0.12 [-0.31, 0.56]	1.75 [-6.37, 11.29]
	Extrinsic-intrinsic	Central	-0.02 [-0.16, 0.10]	0.00 [-0.26, 0.28]	0.02 [-0.34, 0.36]	-0.44 [-6.38, 5.43]
		Peripheral	0.04 [-0.13, 0.23]	0.01 [-0.30, 0.34]	-0.05 [-0.50, 0.38]	-0.55 [-7.33, 6.49]
	Intrinsic-extrinsic	Central	**-0.15 [-0.24, -0.07]**	**-0.32 [-0.51, -0.13]**	**0.47 [0.23, 0.71]**	9.77 [-2.82, 23.46]
		Peripheral	0.02 [-0.08, 0.12]	**-0.22 [-0.43, -0.01]**	0.20 [-0.07, 0.47]	-0.30 [-8.59, 8.74]
2	Intrinsic-intrinsic	Colour cue	**0.42 [0.30, 0.53]**	-0.01 [-0.20, 0.15]	**-0.41 [-0.63, -0.14]**	4.55 [-4.79, 12.63]
		Shape cue	**0.41 [0.29, 0.53]**	0.00 [-0.21, 0.17]	**-0.41 [-0.65, -0.14]**	2.79 [-6.90, 11.13]
		Location cue	**0.46 [0.31, 0.59]**	0.00 [-0.22, 0.23]	**-0.46 [-0.71, -0.18]**	-4.36 [-13.12, 2.93]
	Extrinsic-intrinsic	Colour cue	**0.28 [0.15, 0.41]**	-0.13 [-0.36, 0.06]	-0.15 [-0.42, 0.14]	5.68 [-0.06, 10.74]
		Shape cue	**0.30 [0.17, 0.42]**	-0.08 [-0.33, 0.12]	-0.22 [-0.50, 0.08]	5.85 [-0.17, 11.59]
		Location cue	**0.26 [0.11, 0.42]**	-0.05 [-0.32, 0.20]	-0.21 [-0.53, 0.14]	0.76 [-5.51, 6.28]
	Intrinsic-extrinsic	Colour cue	**0.22 [0.10, 0.35]**	-0.16 [-0.40, 0.06]	-0.06 [-0.35, 0.25]	1.53 [-1.46, 4.50]
		Shape cue	0.08 [-0.03, 0.20]	0.10 [-0.14, 0.34]	-0.19 [-0.48, 0.12]	1.48 [-1.46, 4.22]
		Location cue	**0.57 [0.47, 0.65]**	**-0.26 [-0.47, -0.08]**	**-0.30 [-0.54, -0.03]**	**5.06 [1.98, 7.98]**

Note. Effects in boldface font indicate credible effects.

Furthermore, there were no credible effects of either central (estimated difference = -0.09 [-0.48, 0.32]) or peripheral (estimated difference = -0.09 [-0.44, 0.24]) disruption on binding errors when only the intrinsic features were relevant (i.e., the intrinsic-intrinsic condition). However, given the prior results that central attention disruption had no impact on extrinsic features, we do not interpret these intrinsic feature results too strongly. Overall, we failed to observe support for our hypotheses that disrupting central versus peripheral attention would differently impact binding errors of intrinsic and extrinsic features.

Although our hypotheses focused on binding errors, we also report on target memory (Figure 2A), memory precision (Figure 2B), and guessing (Figure 2D). Disrupting central attention impaired target memory (estimated difference = -0.15 [-0.24, -0.07]) and increased guessing (estimated difference = 0.47 [0.23, 0.71]) relative to a no-disruption baseline in the intrinsic-extrinsic condition. There were no other credible effects of disruption. Thus, the effects of disrupting central attention for recalling an extrinsic feature (i.e., location) appeared to be specific to target memory and guessing rather than binding errors as we had initially predicted.

To be certain that participants did not simply give up during the task, we verified that proportion accuracy on the peripheral (M = 0.95, SD = 0.10) and the central (M = 0.92, SD = 0.05) attention disruption tasks was very high. Coupled with the forthcoming results of Experiment 2, it is thus more likely that randomizing shape and location as additional encoded features alongside colour, as well as intermixing the different probe-target conditions, likely yielded a more challenging task overall for participants. We return to the theoretical implications of this in the General Discussion. Notwithstanding, the current results suggest that disrupting central versus peripheral attention did not impact binding errors as we had predicted, but disrupting central attention may specifically impact target memory and guessing when recalling extrinsic features.

Experiment 2

Method

Materials and Procedure. The materials, procedure, and analysis for Experiment 2 were similar to Experiment 1, except that retro-cues were manipulated during the retention interval (Figure 3). Depending on the block, the retention interval (1250ms) entailed one of four retro-cue conditions. Either the retention interval remained blank with the fixation cross presented at the center of the screen (no-cue condition), or a retro-cue was displayed at the center of the screen for 250ms following the offset of the memory array (500ms) and preceding the onset of the retrieval phase (500ms). The retro-cue indicated with 100% validity which of the items from the memory array would be tested. Specifically, either a circular coloured-dot (colour retro-cue), a dark-grey shape (shape retro-cue), or an arrow pointing to the location of the to-be-tested item (location retro-cue) was presented, depending on the block. There were 50 trials of each retro-cue/probe-target condition.

Design and Data Analysis. The experiment followed a 4 (retro-cue type: none, colour, shape, location) x 3 (probe-target type: intrinsic-intrinsic [colour-shape, shape-colour], extrinsic-intrinsic [location-colour, location-shape], intrinsic-extrinsic [colour-location, shape-location]) within-subjects design. The retro-cue manipulation was blocked and counterbalanced across participants, with the nature of the probe-target type varying randomly within each block. The analytic approach was the same as Experiment 1.

Figure 3

Example Trial Sequence from Experiment 2 that Varied Retro-Cues (None, Colour, Shape, Location) and Probe-Target Type (Intrinsic-Intrinsic, Extrinsic-Intrinsic, Intrinsic-Intrinsic).

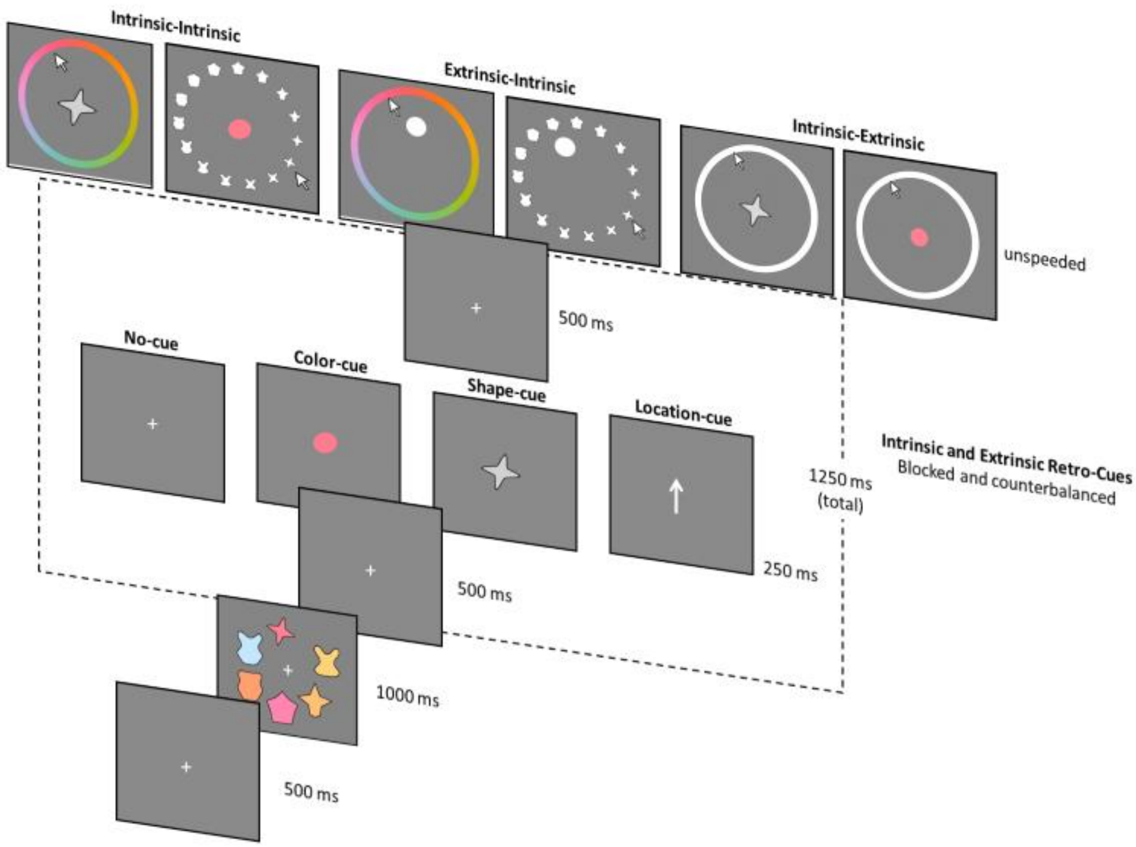

Results and Discussion

Table 3 and Figure 4 summarize our results. We found that extrinsic (i.e., location) retro-cues reduced binding errors in the intrinsic-extrinsic condition compared to the no-cue baseline (estimated difference = -0.26 [-0.47, -0.08]). Although consistent with our hypothesis, in hindsight, we cannot interpret this result given that, in this condition, the retro-cue perfectly matched the target for recall, as was the case in several other conditions (i.e., colour retro-cues probing recall of colour and shape retro-cues probing recall of shape). We return to this issue in the exploratory analyses section. There were no further credible retro-cue effects in binding errors.

Although our hypotheses focused on binding errors, we once again report on target memory (Figure 4A), memory precision (Figure 4B), and guessing (Figure 4D). Overall, we observed credible retro-cue benefits to nearly all target memory parameters except for shape cues in the intrinsic-extrinsic condition (estimated difference = 0.08 [-0.03, 0.20]). Furthermore, all the retro-cues reduced guessing in the intrinsic-intrinsic condition, but Figure 4D shows that guessing was particularly high in this condition's no-cue baseline. Overall, the benefits of retro-cues were most specific to target memory, and most importantly, their efficacy generally did not discriminate between the nature of the retro-cue (whether intrinsic [colour, shape] or extrinsic [location]) nor the nature of the recalled features (intrinsic or extrinsic).

Figure 4

Posterior Parameter Estimates of the Bayesian Hierarchical Mixture Model for the Probability of Target Memory (A), Memory Precision (B), Probability of a Binding Error (C) and Probability of Guessing (D) in Experiment 2. Larger Dark Circles Indicate Group Means, Smaller Faded Circles Indicate Individual Means, and the Error Bars Show the 95% HDIs of the Posterior.

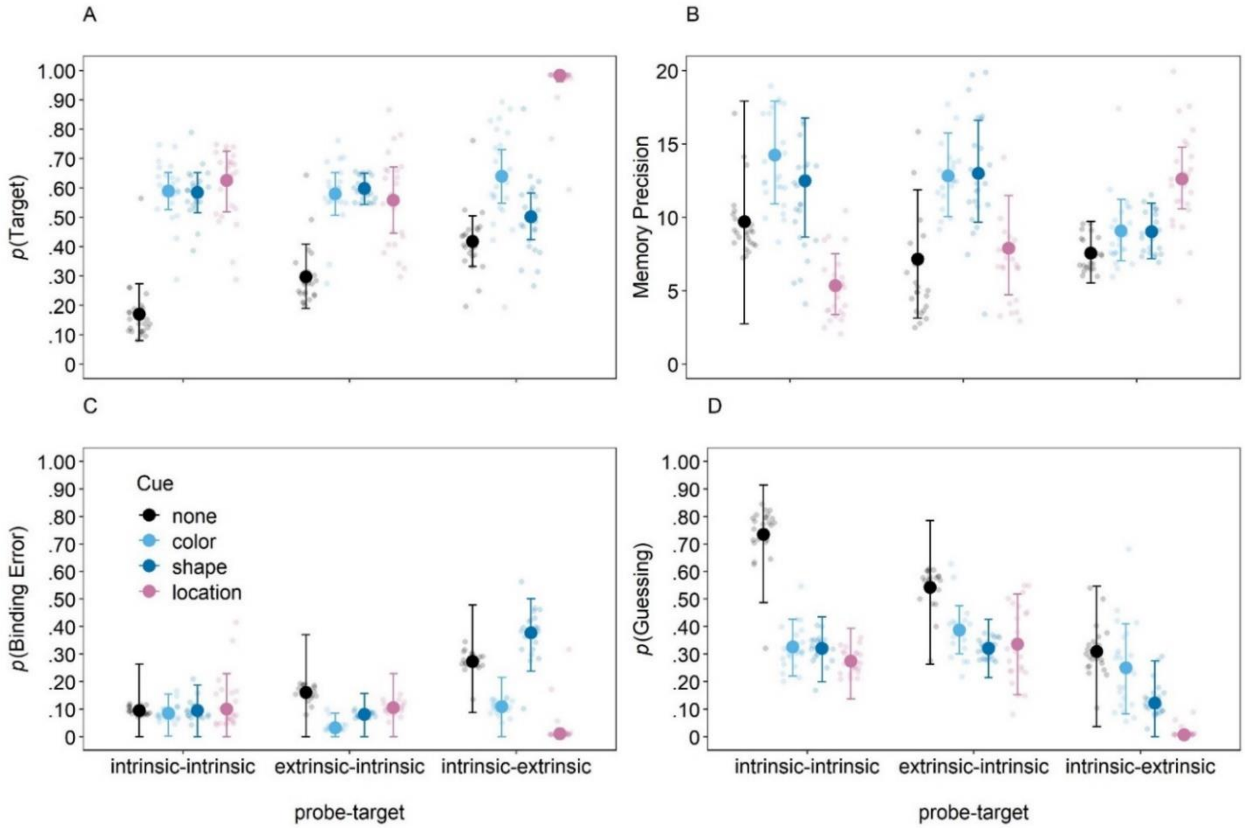

Exploratory Analyses

Given our fully-crossed design, there were trials in which the presented retro-cue matched the to-be-recalled target-feature (e.g., presenting a colour retro-cue and thereafter recalling colour). To understand whether this affected the just-reported confirmatory analysis results, we conducted an additional exploratory analysis that excluded any cue-target matches (i.e., colour-cue/colour-target, shape-cue/shape-target, location-cue/location-target) and then refit the model for each specific probe-target combination (Figure 5 and Table 4).

We observed that location retro-cues were the most consistently beneficial to target memory, except in the location-colour condition (estimated difference = 0.28 [-0.02, 0.59]). Conversely, colour and shape retro-cue benefits were specific to the shape-location condition (estimated differences = 0.35 [0.19, 0.50] and 0.15 [0.01, 0.28], respectively). Thus, location retro-cues were effective largely regardless of whether the probe-target features are intrinsic or extrinsic, whereas intrinsic retro-cues (i.e., shape and colour) only benefitted recalling extrinsic (i.e., location) features. There were few other retro-cue effects observed in the other memory parameters.

These exploratory results clarify those of the confirmatory analyses by suggesting that what at first appear to be largely consistent retro-cue benefits to target memory, regardless of the type of cue or probe-target condition, are in fact largely driven by location retro-cues when recalling colours and shapes, as well as colour and shape retro-cues when recalling a location given a shape probe. This is in line with recent prior work suggesting that intrinsic feature cues, like shape and colour, tend to be less effective overall compared to location retro-cues (Arnicane & Souza, 2021). Thus, although there may be something particularly efficacious about location as an extrinsic feature when guiding attention, our results indicate that this efficacy is largely consistent regardless of whether intrinsic or extrinsic features are relevant to recall.

Figure 5

Posterior Parameter Estimates of the Probability of Target Memory (A), Memory Precision (B), Probability of a Binding Error (C) and Probability of Guessing (D) for each probe-target in Experiment 2. Larger Dark Circles Indicate Group Means, Smaller Faded Circles Indicate Individual Means, and the Error Bars Show the 95% HDIs of the Posterior.

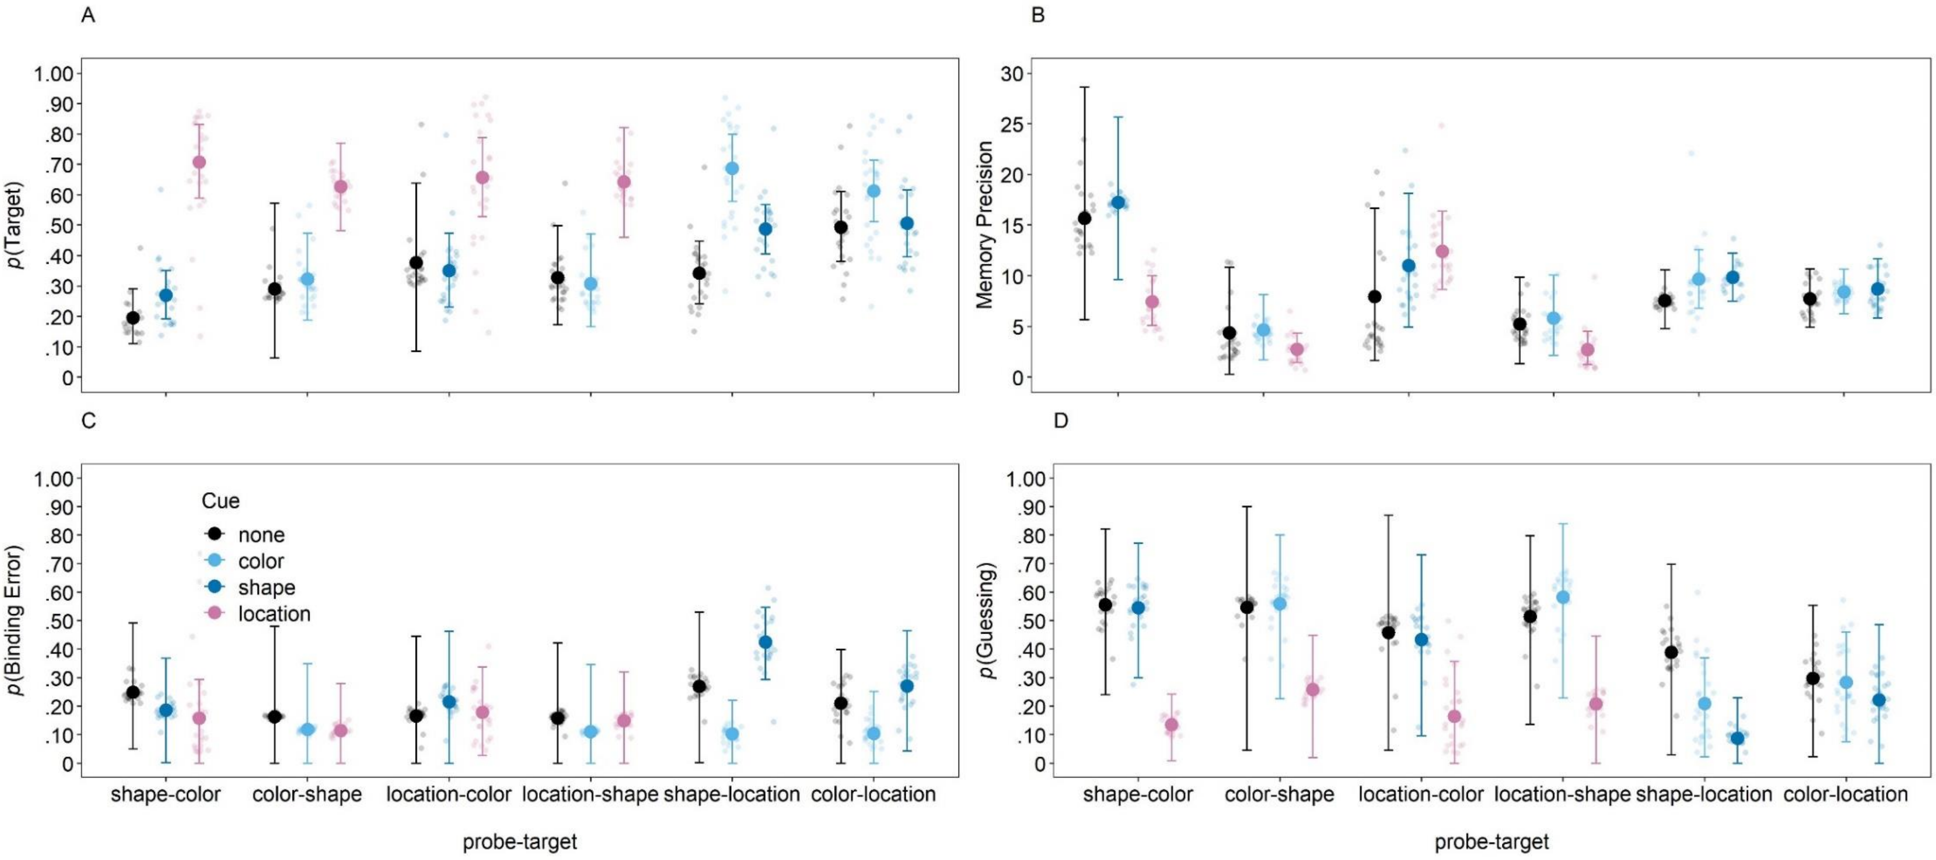

Table 4

Summary of Mean Differences [and 95% HDIs] in Each Memory Parameter for the Exploratory Analyses of Experiment 2.

Probe-Target		Baseline/None vs.	*P*(Target)	*P*(Binding error)	*P*(Guessing)	Precision
Intrinsic-intrinsic	Shape-Colour	Shape	0.07 [-0.05, 0.19]	-0.06 [-0.38, 0.24]	-0.01 [-0.40, 0.39]	1.58 [-14.02, 15.68]
		Location	**0.51 [0.36, 0.67]**	-0.09 [-0.38, 0.17]	**-0.42 [-0.72, -0.09]**	-8.22 [-21.23, 2.42]
	Colour-Shape	Colour	0.03 [-0.29, 0.33]	-0.05 [-0.48, 0.30]	0.01 [-0.48, 0.65]	0.29 [-7.64, 6.85]
		Location	**0.34 [0.02, 0.62]**	-0.05 [-0.43, 0.26]	-0.29 [-0.74, 0.25]	-1.63 [-8.77, 3.38]
Extrinsic-intrinsic	Location-Colour	Shape	-0.03 [-0.33, 0.28]	0.05 [-0.37, 0.42]	-0.02 [-0.56, 0.52]	3.07 [-8.97, 13.80]
		Location	0.28 [-0.02, 0.59]	0.01 [-0.32, 0.30]	-0.29 [-0.76, 0.15]	4.47 [-5.33, 12.92]
	Location-Shape	Colour	-0.02 [-0.26, 0.21]	-0.05 [-0.39, 0.28]	0.07 [-0.40, 0.54]	0.56 [-5.76, 6.93]
		Location	**0.32 [0.06, 0.55]**	-0.01 [-0.35, 0.30]	-0.31 [-0.72, 0.13]	-2.52 [-7.60, 2.01]
Intrinsic-extrinsic	Shape-Location	Colour	**0.35 [0.19, 0.50]**	-0.17 [-0.48, 0.12]	-0.18 [-0.55, 0.22]	2.14 [-1.95, 6.40]
		Shape	**0.15 [0.01, 0.28]**	0.16 [-0.15, 0.45]	-0.30 [-0.66, 0.07]	2.32 [-1.57, 6.02]
	Colour-Location	Colour	0.12 [-0.04, 0.27]	-0.11 [-0.37, 0.16]	-0.01 [-0.36, 0.31]	0.69 [-3.09, 4.32]
		Shape	0.01 [-0.14, 0.18]	0.06 [-0.26, 0.38]	-0.08 [-0.48, 0.33]	0.97 [-3.35, 5.06]

Note. Effects in boldface font indicate credible effects.

General Discussion

Overall, the results of the current experiments suggest that attention does not discriminate between different types of feature bindings in visual WM as we had predicted. In Experiment 1, we showed that disrupting central and peripheral attention did not respectively increase extrinsic (i.e., location) and intrinsic (i.e., colour, shape) binding errors relative to a no-disruption baseline. Instead, disrupting central attention only reduced target memory and increased guessing when recalling the extrinsic feature. Further against the predicted dissociation of extrinsic and intrinsic binding errors, Experiment 2 showed that extrinsic retro-cues were the most effective to increase target memory, regardless of whether the probe-target features were intrinsic or extrinsic. These results thus suggest that manipulating attention impacts target recall in visual WM, with no consistent distinction between intrinsic and extrinsic features in the continuous reproduction paradigm used here.

Before any strong interpretation of the results, it is prudent to compare them first to that of previous similar work, such as Souza and Oberauer (2017), whose Experiments 1A and 1B inspired the design of Experiment 1 and Souza (2016) which inspired the design of Experiment 2. The results of those experiments showed much greater target memory in the no-disruption baseline conditions (Experiment 1A: estimate = 0.54 [0.47, 0.61]; Experiment 1B: estimate = 0.48 [0.39, 0.56]; see OSF for details) compared to the same relative extrinsic-intrinsic condition of our Experiment 1 (estimate = 0.19 [0.11, 0.28]). Target memory was also much lower overall in our Experiment 2 (no-cue location-colour estimate = 0.38 [0.09, 0.64]; retro-cue location-colour estimate = 0.66 [0.53, 0.79]) compared to the same relative conditions of Souza (2016; no-cue location-colour estimate = 0.63 [0.54, 0.72]; retro-cue location-colour estimate = 0.87[0.78, 0.95]). It may be that the novelty of the current design, which randomized three different features of colour, shape, and location within the same memory array while randomly intermixing different probe-target conditions increased the overall difficulty of the task relative to similar prior work. This may not be a mere methodological difference given earlier seminal work showing that increasing the number of features to encode and maintain does not impair change detection (e.g., Luck & Vogel, 1997). The current paradigm of continuous reproduction of randomly probed features will thus be useful to future work in that it may better reveal the increased demand of additional features in visual WM. Furthermore, adapting the paradigm so that its overall difficulty does not overwhelm participants will be important, for example, by calibrating the presented set size of the arrays to each individual participant's ability level (e.g., Loaiza & Souza, 2019).

A further caveat to the pattern of results is that binding errors were relatively low across conditions of both experiments. Furthermore, binding errors tended to occur more frequently in the intrinsic-extrinsic baseline condition than the other baseline conditions in both experiments. This may suggest that locations as targets yield more binding errors than shape and colour in baseline conditions, which will require further investigation. It is important to note that this pattern does not impact the current analyses or conclusions given that the critical comparisons for the research questions were conducted by comparing the effect of disrupting or guiding attention to the relevant baseline within each probe-target combination. Furthermore, the general low rate of binding errors makes it difficult to determine whether there were no true effects of our conditions on binding errors or simply a reduced opportunity to observe any effects, such that binding errors may have been too low overall to be sensitive to disruption or retro-cue effects.

Notwithstanding, the pattern of results for target memory makes it unlikely that recalling intrinsic versus extrinsic features varies depending on attentional disruption or retro-cues as we had predicted. Although disrupting central attention hampered target recall of the extrinsic feature in Experiment 1, we observed retro-cue effects in Experiment 2 that did not consistently distinguish between extrinsic and intrinsic features. Extrinsic retro-cues most effectively enhanced target memory compared to other cues, consistent with recent work (Arnicane & Souza, 2021), but in one case intrinsic retro-cues also enhanced target memory of extrinsic features. Thus, we interpret these results to suggest that manipulating attention does not differently impact intrinsic and extrinsic feature binding, calling into question whether a distinction should be made between them. However, further work is required to replicate this pattern of results in other paradigms to determine whether there truly are no distinctions between intrinsic and extrinsic feature binding.

Chapter 3

Reducing set size did not reveal distinctions between central and peripheral disruption.

Rationale

Prior research has suggested that visual features of an item can be classified as intrinsic or extrinsic (see Chapter 2; also see Ecker et al., 2013). The results of Experiment 1 in Chapter 2 showed that there was no detrimental impact of attention disruption on binding errors, regardless of the type of disruption (central or peripheral) or type of feature (intrinsic or extrinsic). However, the central attention disruption reduced target recall and increased guessing relative to a no-disruption baseline. The question that is addressed in this chapter is whether a relatively high number of items to remember (i.e., a set size of 6) limited the opportunity to observe more substantial impacts of the attention disruption manipulation.

Throughout the WM literature it has been argued that WM capacity (i.e., one's ability to keep multiple items active for a short period of time) may be limited by a fixed number of items (e.g., Cowan, 2001; see Luck & Vogel, 2013 for a review). For example, across a series of experiments, Luck and Vogel (1997) demonstrated that processing and retaining visual information occurs in an object-based manner, such that representations of an item are retained as feature bindings rather than as individual features. In their experiments, the researchers manipulated the set size of to-be-retained items (ranging from 1-12 items) and the feature type (i.e., colour, orientation, or both) relevant for change detection. They found that performance systematically declined after 4 items were tested. Luck and Vogel (1997) argued that visual WM retains a limited number of integrated item representations that can each comprise many individual features. Thus, the quantity of individual features that can be retained is limited by the number of object representations that are retained. In agreement, Vogel and colleagues (2001) argued that visual WM is limited by the number of items that are to-be-retained rather than the number of features that make up an item representation. Thus, the number of to-be-remembered items may limit WM maintenance, regardless of the number of features that can be retained.

One account of visual information storage in WM has argued that item representations may be stored in limited capacity slots and require additional resources to determine which item representations are goal relevant (i.e., slots + resource hypothesis; Zhang & Luck, 2008). Moreover, Zhang and Luck (2008) argued that high-quality representations of a single item occur by distributing individual features across the limited capacity slots (i.e., slots + averaging hypothesis). This suggests that although WM is limited by the number of items that can be retained, one can improve the quality of a to-be-retained goal-relevant item by selectively maintaining the features of that item. More recent research has demonstrated that this may be the case. For example, Oberauer and Eichenberger (2013) demonstrated reduced change detection accuracy while retaining items with multiple features compared to a single-featured items. This highlights that WM capacity may not only be limited by the number of objects that can be retained as Luck and Vogel (1997) suggested, but it may also be limited by the number of features that can be retained. Thus, it is important to understand whether the results of Experiment 1 in Chapter 2 may be attributed to exceeding participants' WM capacity while retaining 6 multi-feature items (composed of: 6 colour features + 6 shape features + 6 location features). If this is the case, then reducing the number of to-be-retained items should allow for a greater understanding of the initial hypotheses in Experiment 1 of Chapter 2.

Current Study

The current study therefore aimed to replicate Experiment 1 of Chapter 2 but reducing the number of to-be-remembered items from 6 to 4. The set size was reduced to 4, rather than

1 or 2 items, due to prior work suggesting that WM can hold up-to 4 items active at once (Cowan, 2001; Luck & Vogel, 1997; Zhang & Luck, 2008). Thus, the reduced set size should offer greater task performance in the current study, compared to the greater set size of 6 in Experiment 1 of Chapter 2. The hypotheses were the same as before (see Table 1 for full summary): If domain-specific peripheral attention is required to selectively maintain intrinsic features, then disrupting peripheral attention should selectively increase intrinsic binding errors but not extrinsic binding errors, relative to a no-disruption baseline. If central attention selectively maintains extrinsic bindings, then disrupting central attention should selectively increase extrinsic binding errors but not intrinsic binding errors relative to a no-disruption baseline. Reducing the set size from 6 items to 4 items should allow for a greater understanding of disrupting attention on intrinsic and extrinsic binding errors as the reduced set size is within the WM capacity. Therefore it can be argued that with the reduced set size in the current study, and while peripheral and central attention are disrupted, the intrinsic and extrinsic binding errors should be reflective of the disruption-type, rather than due to exceeding WM capacity (i.e., the higher set size of Experiment 1 of Chapter 2), or possible noise.

Method

Participants. As in Experiment 1 of Chapter 2, 24 complete and valid responses were collected from participants aged 18-35 through the University of Essex subject pool and Prolific, by voluntary and paid participation. Participants were awarded either partial course credit, or monetary compensation (UK £7.50 per hour). All participants reported normal or corrected-to-normal vision, colour vision and hearing. Prior to starting the experiment, participants undertook a colour blindness task to ensure their colour vision would not hinder their performance. Forty-one participants were excluded from the current study due to not passing phase 1 (i.e., achieving 100% on a colour blindness test, and 80% on a visual and an auditory task). In line with Experiment 1 of Chapter 2, participants were offered multiple opportunities to practice the tasks before completing the test that determined whether they passed.

Design, Materials, Procedure, and Analysis. The design, materials, procedure, and analysis were identical to Chapter 2 Experiment 1, with the only exception being that the set size was reduced from six to four coloured shapes to remember.[9] Further Bayesian analyses were conducted to understand whether reducing the set size in the current chapter impacted recall performance. A Bayesian approach allows for a comparison between two models based on the Bayes Factor (BF) statistic, which conveys the ratio between the likelihoods of the data under the two compared models. For example, to determine the evidence for a main effect of set size, one would compare the likelihood of a model assuming an effect (M_1) relative to the null model (M_0; BF_{10}). BFs are interpreted continuously, such that, a BF between 1 and 3 (or 0.01 to 0.03) is interpreted as ambiguous, and BFs greater than 3 and 10 (0.03 and 0.10) are substantial and strong evidence for/against the model in the numerator relative to the model in the denominator. The anonymised data files, analysis scripts and experiment materials are available on the open science framework (OSF; https://tinyurl.com/2ukk6anx).

Results

Against my predictions, there were no credible effects of the attention-disruption task in the binding error parameter estimates, irrespective of the probe-target type (see Table 5). Thus, these results suggest that disrupting central and peripheral attention does not increase intrinsic and extrinsic binding errors relative to the baseline condition, regardless of the change in set size between experiments. Upon closer inspection (and as can be seen in Figure 6C), it

[9] I observed an unknown error wherein 100 recorded responses (i.e., 0.47% of the data) were greater than 360 degrees (i.e., the maximum possible response). These responses were omitted from analysis.

is likely that the overall low binding errors may have constrained the opportunity to observe any effect, as in Chapter 2. Contrary to Experiment 1 of Chapter 2, no credible effects of attention disruption were observed in target memory or guessing, irrespective of the nature of the probe (see Table 5). This suggests that the credible differences in these parameters in the intrinsic-extrinsic condition of Experiment 1 of Chapter 2 may be attributed to set sizes greater than 4. Therefore, central attention may be required to accurately recall extrinsic (location) features, particularly when more than 4 items are to-be-retained.

An alternative explanation for the observed effects in Experiment 1 of Chapter 2 may be that rather than relying on central attention to accurately recall extrinsic features while WM capacity has exceeded, WM performance is reduced and therefore the effects observed in Experiment 1 of Chapter 2 may be attributed to noise rather than disruption effects. Upon closer inspection of Figure 6A, despite the peripheral and central attention disruption conditions, target memory is moderate and high. Thus suggesting that accurately recalling intrinsic and extrinsic bindings may rely on one’s WM capacity. Therefore, further exploratory analyses of task performance (i.e., recall error) between the current study and Experiment 1 of Chapter 2 is required to indicate whether task performance improved with the reduced set size of 4 items in the current study, in comparison to the higher set size of 6 items in Experiment 1 in Chapter 2.

Table 5

Summary of Mean Differences [and 95% HDIs] in Each Memory Parameter.

Probe-target type	Baseline/None vs.	*P*(Target)	*P*(Binding Error)	*P*(Guessing)	Precision
Intrinsic-Intrinsic	Central	0.38 [-0.02, 0.75]	-0.02 [-0.18, 0.11]	-0.36 [-0.78, 0.10]	-4.81 [-11.43, 2.66]
	Peripheral	0.20 [-0.21, 0.58]	0.05 [-0.12, 0.26]	-0.25 [-0.67, 0.23]	-1.94 [-9.20, 5.69]
Extrinsic-Intrinsic	Central	0.13 [-0.24, 0.46]	-0.01 [-0.09, 0.07]	-0.13 [-0.48, 0.27]	-4.07 [-10.32, 2.71]
	Peripheral	0.11 [-0.04, 0.27]	0.11 [-0.04, 0.27]	-0.22 [-0.50, 0.07]	-0.46 [-6.75, 6.35]
Intrinsic-Extrinsic	Central	-0.07 [-0.33, 0.18]	0.01 [-0.21, 0.23]	0.06 [-0.25, 0.44]	-1.82 [-5.36, 1.72]
	Peripheral	-0.05 [-0.24, 0.15]	-0.05 [-0.24, 0.15]	0.10 [-0.15, 0.34]	-0.07 [-2.44, 2.32]

Figure 6

Posterior Parameter Estimates for the Probability of Target Memory (A), Memory Precision (B), Binding Errors (C) and Probability of Guessing. Larger Dark Circles Indicate Group Means, Smaller Faded Circles Indicate Individual Means, and The Error Bars Show the 95% HDIs of the Posterior.

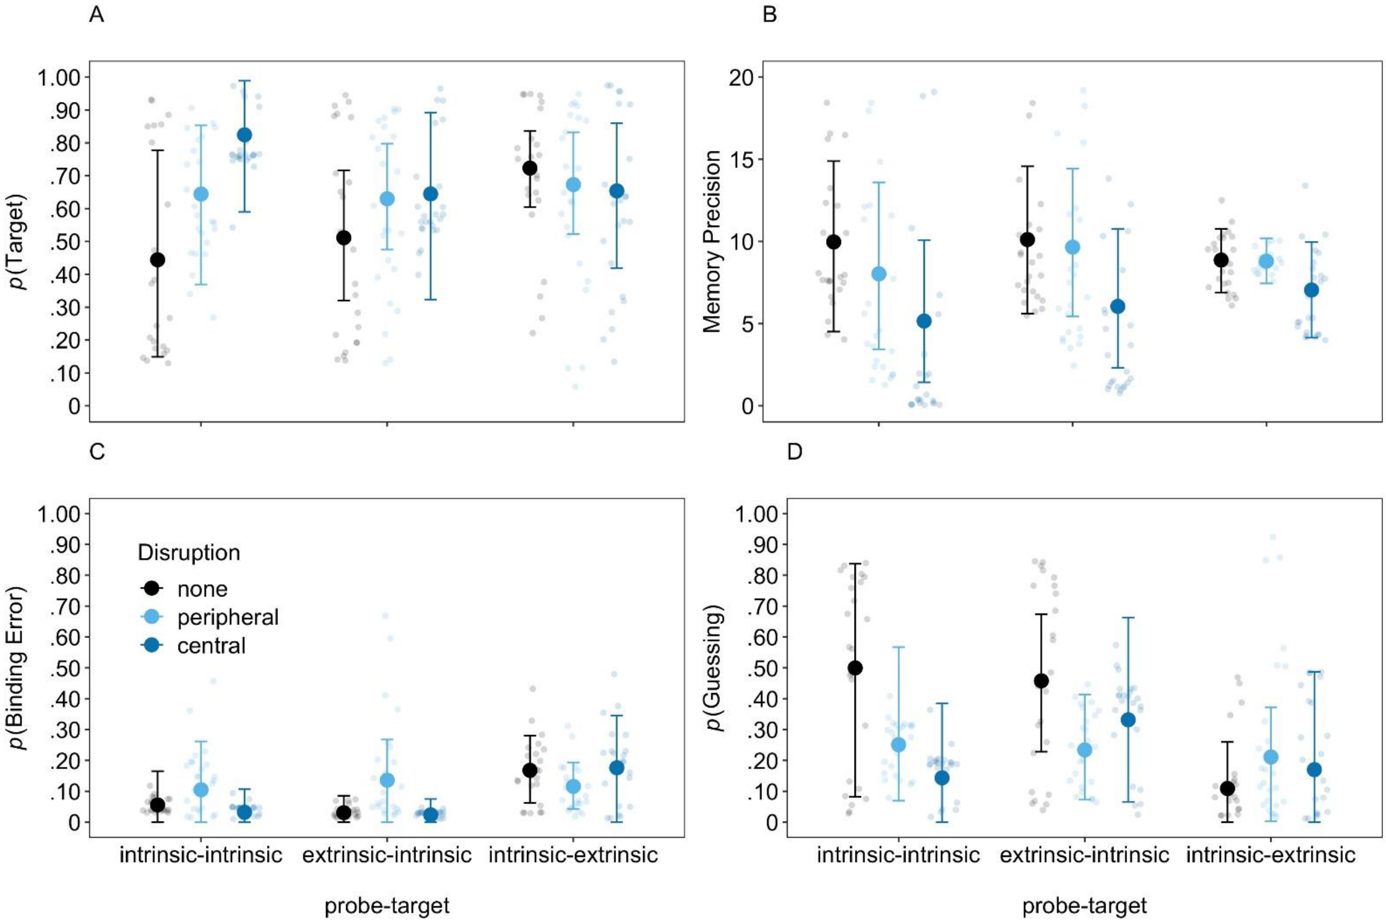

Exploratory Analyses

To determine whether reducing the number of to-be-retained items improved overall task performance in the current study, in comparison to the higher set size in Experiment 1 of Chapter 2 (see Figure 7 for the plotted distribution of recall errors), a 3 (intrinsic-intrinsic, extrinsic-intrinsic, intrinsic-extrinsic) x 3 (none, peripheral, central) x 2 (4, 6) mixed BANOVA was conducted using the R package BayesFactor (Morey et al., 2022). The model including all three main effects was the best model, $BF_{10} = 1.34 \times 10^{31}$. The model including all three main effects was ambiguously preferred over the next-best model including the interaction of set size and binding-type, and binding-type and disruption-type (BF = 2.73). Table 6 shows paired Bayesian t-tests comparing the recall error of set size 4 (current chapter) and 6 (Experiment 2 of Chapter 2) for each of the binding types in their respective disruption-type, against the set sizes. These tests indicated very strong evidence that reducing the set size in the current chapter improved recall performance, irrespective of binding- or disruption-type. Taken together, these findings suggest that reducing the set size in the current study improved overall recall performance.

Figure 7

Distribution of Recall Errors per Disruption-Type and Set Size for Intrinsic-Intrinsic, Extrinsic-Intrinsic, and Intrinsic-Extrinsic Probe-Target Types.

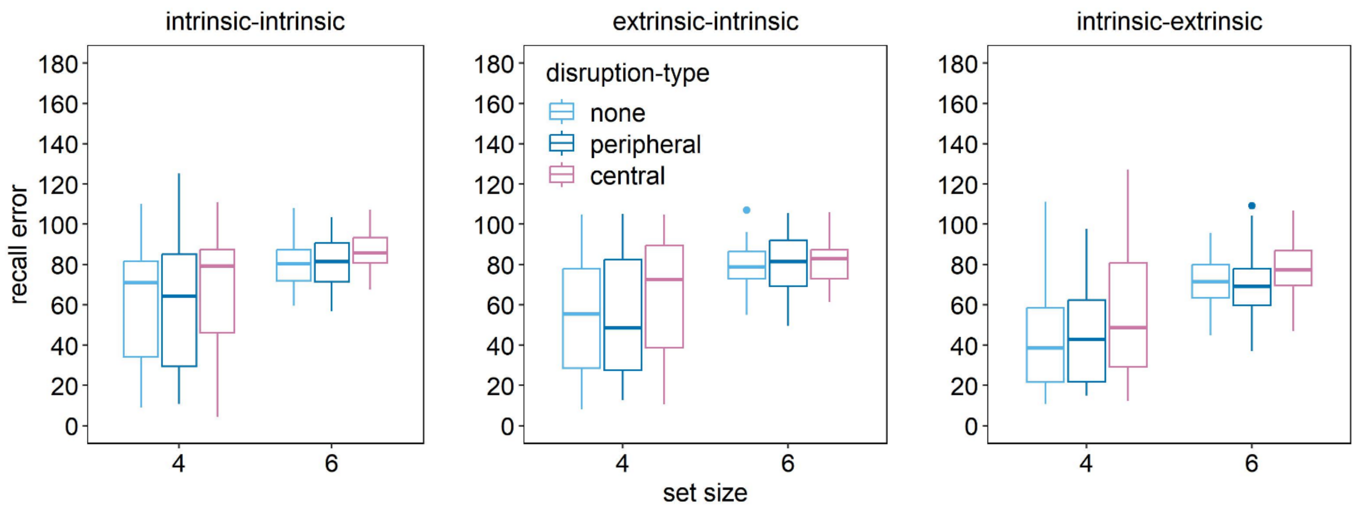

Table 6

Bayesian T-Test Comparing Recall Error Between Set Sizes 4 and 6 of Each Binding Type and Disruption Condition.

	Disruption Condition		
Binding Type	None	Peripheral	Central
Intrinsic-Intrinsic	457.87	489.17	821.82
Extrinsic-Intrinsic	7636.74	157918	172.82
Intrinsic-Extrinsic	9503945	230300	46271.62

Discussion

To understand whether the findings of Chapter 2, Experiment 1 were due to exceeding WM capacity, the paradigm was conducted again, except the number of to-be-remembered items was reduced from 6 to 4. The findings suggested that reducing the number to-be-remembered items did not change the null effect of attention disruption on binding errors. Furthermore, no credible effects were observed in the remaining model parameters (target memory, precision, and guessing) irrespective of the nature of the probe-target combination. Moreover, additional analyses comparing recall error between Experiment 1 in Chapter 2 and the current study revealed very strong evidence in favour of task performance improving with the reduced set size. Thus, as has been observed many times in the literature, WM capacity is limited, such that memory performance improves when fewer items are to be kept active in WM (Cowan, 2001; Cowan et al., 2013; Zhang & Luck, 2008).

Given the design of these experiments, one may ask whether the improved performance occurred due to reducing the number of objects (from 6 to 4) or the number of to-be-remembered features (from 18 to 12). It is not possible to know this in the current work given that both the number of objects and features were reduced simultaneously. Further research is required to determine this question. For example, one may investigate whether features or objects contribute to the capacity limit by adapting the current paradigm calibrate the number of features presented on each trial (i.e., increasing or decreasing the number of features presented) based on mean recall performance on the last *n* trials. Alternatively, calibrating the number of objects that are presented on each trial (using the same aforementioned approach, but the number of features per object remain constant on each trial) could also determine whether one's WM capacity depends on one's ability to maintain multiple objects, features, or both. For example, if adapting the number of features per object impacts recall performance more strongly than the number of objects, then this would be evidence for a capacity limit of features rather than objects. Thus, this approach may offer a novel insight into the current debate as to whether items are represented and stored as bound representations or as individual features (see Luck & Vogel, 2013; also see Schneegans & Bays, 2019 for a review), as one would be able to determine whether recall performance is hindered by the increasing number of presented feature types.

To summarize, the study of this chapter examined whether reducing the set size would determine whether disrupting two different attentional resources would increase intrinsic and extrinsic binding errors. While I did not find that disrupting peripheral and central attention increased intrinsic and extrinsic binding errors respectively, my findings suggest that task performance was improved with the reduced set size in the current study. As I explained previously, the paradigm may be useful for distinguishing whether the number of objects versus their individual features may contribute to set size effects commonly observed in the WM literature.

Chapter 4

Guiding attention via retro-cues improves target memory and guessing but does not impact binding errors.

Rationale

A longstanding theoretical debate in visual working memory (WM) is whether visual objects are encoded, stored, and maintained in an object-based manner (i.e., features of an object are bound together and the bound representations are maintained), or in a feature-based manner (i.e., features of an item are encoded, stored, and maintained independently; for a review, see Luck & Vogel, 2013; also see Schneegans & Bays, 2019). One way in which this has been investigated is by considering the role of attention in visual WM. Congruent with the object-based view, Allen and colleagues (2006) suggested that attention is not required for feature binding, rather, bound representations are processed automatically. Incongruent with the object-based view, Treisman and Gelade (1980) argued that the processing of multi-feature objects (i.e., objects composed of multiple features such as colour, shape, and spatial-location) relies on features of an item being registered early in the visual field, but attention binds relevant features together into a cohesive representation. The researchers argued that feature binding is driven by spatial-location features, and as such non-location features (i.e., colour, or shape) are bound to their respective location. This suggests that features of an item may be maintained independently. Further work has argued that location features may be privileged in visual WM, relative to non-location features (Rajsic & Wilson, 2014; Schneegans & Bays, 2017). This raises the question as to whether locations are privileged in visual WM, and whether attention selectively prioritises locations compared to non-location features. As such, are objects represented in an object-based manner without the requirement of attention, or are objects represented in a feature-based manner in visual WM?

Seminal research investigating visual WM capacity limits has suggested that objects are represented as bound representations (Luck & Vogel, 1997). Luck and Vogel (1997) demonstrated that in a visual WM task in which participants were required to maintain individual features of an item (i.e., colour, orientation, size, and gap) alongside the bound representation (i.e., the whole object), participants were as accurate at reporting on the bound representation at test as they were at reporting individual features. The similarity in accuracy led the researchers to suggest that objects are maintained as bound representations rather than as individual features, despite whole objects being composed of multiple individual features. Further work has demonstrated that visual WM is limited by the number of bound representations it can maintain (Vogel et al., 2001). Across 16 Experiments, Vogel and colleagues (2001) argued that visual WM is constrained by the number of bound representations it can hold active, rather than the number of individual features it can maintain. Moreover, the researchers argued that the maintenance of bound representations does not rely on attention. Consistent with the view that attention is not required for the maintenance of bound representations, in a series of experiments Allen and colleagues (2006) observed that disrupting attention by means of a secondary task (such as counting backwards) did not reduce recall accuracy of bound representations to a greater extent than the individual features themselves. Thus, the object-based view suggests that attentional resources are not required during visual WM maintenance of multi-feature objects that are represented in an object-based manner, with no limitation on the number of features that are bound within the representation. Alternative research has argued that visual information may not be represented in an object-based manner, rather features of an item are represented cohesively in visual WM (Foungie et al., 2013). By manipulating whether participants were required to retain 5 coloured and orientated triangles, or 10 separate features (i.e., 5 colours and 5 orientations) simultaneously, Foungie and colleagues (2013) observed that participants were able to correctly recall item features when fewer items were to-be-retained. Foungie and colleagues argued that the higher recall performance in the 5-object condition may be due to there being fewer opportunities for participants to guess during retrieval due to being able to bind the colour and orientations into cohesive object files. Arguably, this suggests that an object-based limit as proposed by earlier work (e.g., Luck & Vogel, 1997; Vogel et al., 2001) may reflect one's ability to process and bind features into a cohesive representation. A fundamental framework from perception research, the feature integration theory (FIT; Treisman & Gelade, 1980) proposed that features of an item are registered in the visual field independently, but maintaining individual features relies on internal attention toward their shared location, thus enabling interdependent binding between non-location and relevant location features. While FIT is considered a framework in the immediate perception and processing of visual information, it has been noted that FIT is an integral hypothesis for the way in which visual features are bound together into object files in WM (Wolfe, 2020). Following on Vogel and colleagues' (2001) experiments, Wheeler and Treisman (2002) observed that visual WM is not limited by the number of bound representations, but rather, it is limited by the number of individual features it can hold active that are composed of the same feature type (i.e., holding multiple colours active). Furthermore, the researchers proposed that individual features are stored in their respective feature dimensions (i.e., colours are stored in a colour-specific store, and locations are stored in a location-specific store). This therefore suggests that visual WM may hold a limited number of features active in their respective feature dimension store.

If features are indeed maintained individually but linked to a shared location, then this implies that factors such as attention may impact some features (i.e., location) and not others. Rajsic and Wilson (2014; Experiment 1) addressed this question using a continuous report paradigm to distinguish between to-be-recalled features (colour or location). The researchers found that participants were more likely to correctly recall a location feature while probed with its associated colour, in contrast to recalling a colour feature while probed with a location feature. Moreover, the researchers found that swap errors (i.e., recalling presented non-targets) occurred more often when recalling locations than colours. These findings suggest that there may be differences between the processing of non-location and location-based features, such that non-location features may be "mapped" onto (or coded in conjunction with) their respective location features (see Johnson & Pashler, 1990). Rajsic and Wilson's (2014) findings suggests that features are maintained in an interdependent fashion wherein location features are privileged. In a similar vein, Schneegans and Bays (2017) further suggested that spatial attention plays a critical role in the binding of location and non-location features. In a continuous-report paradigm, in which participants' attention was guided via a retro-cue during the retention interval (see Souza & Oberauer, 2016 for a review of retro-cue paradigms in visual WM), participants were required to recreate the associated feature (Schneegans & Bays, 2017). In Experiment 1, Schneegans and Bays (2017) observed that while a colour cue was presented, participants were more prone to location swap errors than when a location cue was presented. In Experiment 2, the researchers guided attention via non-location retro-cues (colour and orientation) and participants were required to report the corresponding orientation and location, or colour and location features sequentially using their respective feature wheels. Feature retrieval was balanced across blocks, in that the order of feature retrieval (i.e., colour first, location second, versus location first, colour second) varied.

Schneegans and Bays observed that the error standard deviations were lower during location retrieval, compared to the error standard deviations of the non-location features. Moreover, the researchers did not observe a significant effect of the retrieval order. Taking these findings together, this therefore suggests that guiding attention towards an in-memory location by means of a retro-cue improves the quality of a non-location representation. Furthermore, guiding attention via a non-location retro-cue has little impact on location response errors, but retrieving a colour/orientation feature when cued with an orientation/colour increases response errors. This highlights that location features may be privileged in visual WM and attention may be selectively orientated towards locations. Therefore, the goal of the current study was to further investigate whether locations are indeed privileged during WM maintenance and whether guiding attention towards in-memory features selectively benefits location.

Current Study

The current study adapted Experiment 2 in Chapter 2's paradigm to investigate whether location features are privileged compared to non-location features (in this case, colour), such that guiding attention towards an in-memory location feature (via a retro-cue) particularly benefits the prioritisation of an associated feature. In addition, to understand whether location features are privileged in WM while attention is not guided via a retro-cue, the current study compared colour and location retrieval during the baseline (no-cue) condition. The current study adapted Experiment 2 of Chapter 2 by only including the baseline, colour, and location retro-cues during the retention interval. The current study disentangled feature binding types by manipulating the probe-target combination during retrieval (i.e. probe-target combinations of colour-location or location-colour) to examine whether there are differences between the maintenance of colour and location features. In comparison to Experiment 2 of Chapter 2, rather than intermixing the probe-target types during retrieval in the retro-cue conditions, in the current study the probe-target type was matched to the retro-cue condition. For example, while a colour retro-cue was presented, participants were always required to retrieve the corresponding location feature (see Materials and Procedure for further details). Multi-feature objects were presented during encoding, given prior work suggesting that the number of features an item comprises does not impact visual WM capacity (Luck & Vogel, 1997; Luria & Vogel, 2011; Vogel et al., 2001). A hierarchical Bayesian three-parameter mixture model was applied to observed recall error (Oberauer et al., 2017), with particular interest in binding errors given Rajsic and Wilson's (2014; Experiment 1) finding that swap errors were higher while recalling location compared to colour, consistent with work suggesting that location features may be privileged in WM (e.g., Johnson & Pashler, 1990; Schneegans & Bays, 2017; Treisman & Gelade, 1980). Firstly, it was hypothesised that if location features are privileged without the requirement of selective attentional resources, then binding errors should be lower when recalling a location compared to recalling colour in the no-cue/baseline condition. Secondly, if attention impacts the binding of location and non-location features, then guiding attention via a location-based retro-cue should reduce colour binding errors relative to a no-cue/baseline condition. Conversely, guiding attention via a colour retro-cue should not reduce binding errors when recalling a location feature relative to a no-cue baseline condition. Finally, if non-location features are not bound to their respective locations and attentional resources do not selectively maintain location features, then binding errors should be similar, irrespective of the cue condition (i.e., none, colour-cue, and location-cue).

Method

Participants. In line with Experiment 2 in Chapter 2, 24 complete and valid datasets were collected from participants aged 18-35 through Prolific (Male = 11, Female = 13; M_{age} = 25.16).

Participants received monetary compensation (UK £7.50 per hour) for their participation. All participants reported normal or corrected-to-normal vision and colour vision and did not participate in the experiments of the previous chapter. Furthermore, eight participants were excluded for not passing a colour blindness task, and a further 8 participants were excluded due to not completing the whole study.

Materials and Procedure. The experiment was conducted online through lab.js (Henninger et al., 2022) hosted on the JATOS server Mindprobe (Lange et al., 2015). After a brief demographic questionnaire and colour blindness task, participants completed three blocks (100 trials each) of a visual WM task, with 6 practice trials proceeding each block. Each trial began with a fixation cross presented for 500 ms followed by an array of six coloured shapes simultaneously presented in random locations around an invisible circle for 1000 ms. The colour, shape, and location features of each item were randomly sampled from 360^{o} of continuous values, with a minimum distance of 60^{o} of separation in each feature dimension from the other memoranda in the array: The colours were sampled along a circle in the CIELAB colour space (with L = 70, a = 20, b = 38 , and radius = 60), the shapes were randomly drawn from a shape wheel (Li et al., 2020), and the locations were randomly drawn along an invisible circle (radius = 150^{o}). Depending on the block, the retention interval (1250 ms) entailed one of three retro-cue conditions (see Figure 8). The retro-cue conditions (none, colour, location) were blocked and counterbalanced, and followed a within-subjects design. Following the offset of the memory array (500 ms), the retention interval either remained blank with the fixation cross presented at the centre of the screen (none), or a retro-cue was displayed at the centre of the screen for 250 ms following the offset of the memory array (500 ms) and preceding the onset of the retrieval phase (500 ms). The retro-cue indicated with 100% validity which of the items from the memory array would be tested.

Specifically, either a circular coloured-dot (colour retro-cue), or an arrow pointing to the location of the to-be-tested item (location retro-cue) was presented, depending on the block.

During the retrieval phase, participants were probed with either the colour or location of one item of the memory array to prompt recall of item's location or colour, respectively. The nature of the probe-type (colour or location) depended on the block: During the no-cue block, there was a random mix of colour and location probes presented at the centre of the screen. During the colour-cue block, participants always received a colour-probe (presented at the centre of the screen) to recall the item's location, whereas during the location-cue block, participants always received a location-probe to recall the item's colour. For each probe-target-type, retrieval occurred along a continuous location or colour wheel (depending on the to-be-recalled target feature), wherein participants used a mouse to click along the corresponding wheel. The recall attempt was unspeeded, and participants were instructed to prioritize accuracy over speed in their responses. There was an inter-trial interval of 1000 ms followed by a screen that said "Ready?" – to which participants pressed the spacebar to proceed to the next trial. After the practice trials and every 10 test trials, participants received feedback about their average recall accuracy (expressed as a percentage of their mean reproduction error, i.e., 100 – 100 * mean error/180). Participants were offered a break after every 50 trials.

Figure 8

Example Trial Sequence that Varied Retro-Cues (Baseline/None, Colour, Location) and Probe-Target Types (Colour-Location or Location-Colour, Colour-Location, Location-Colour).

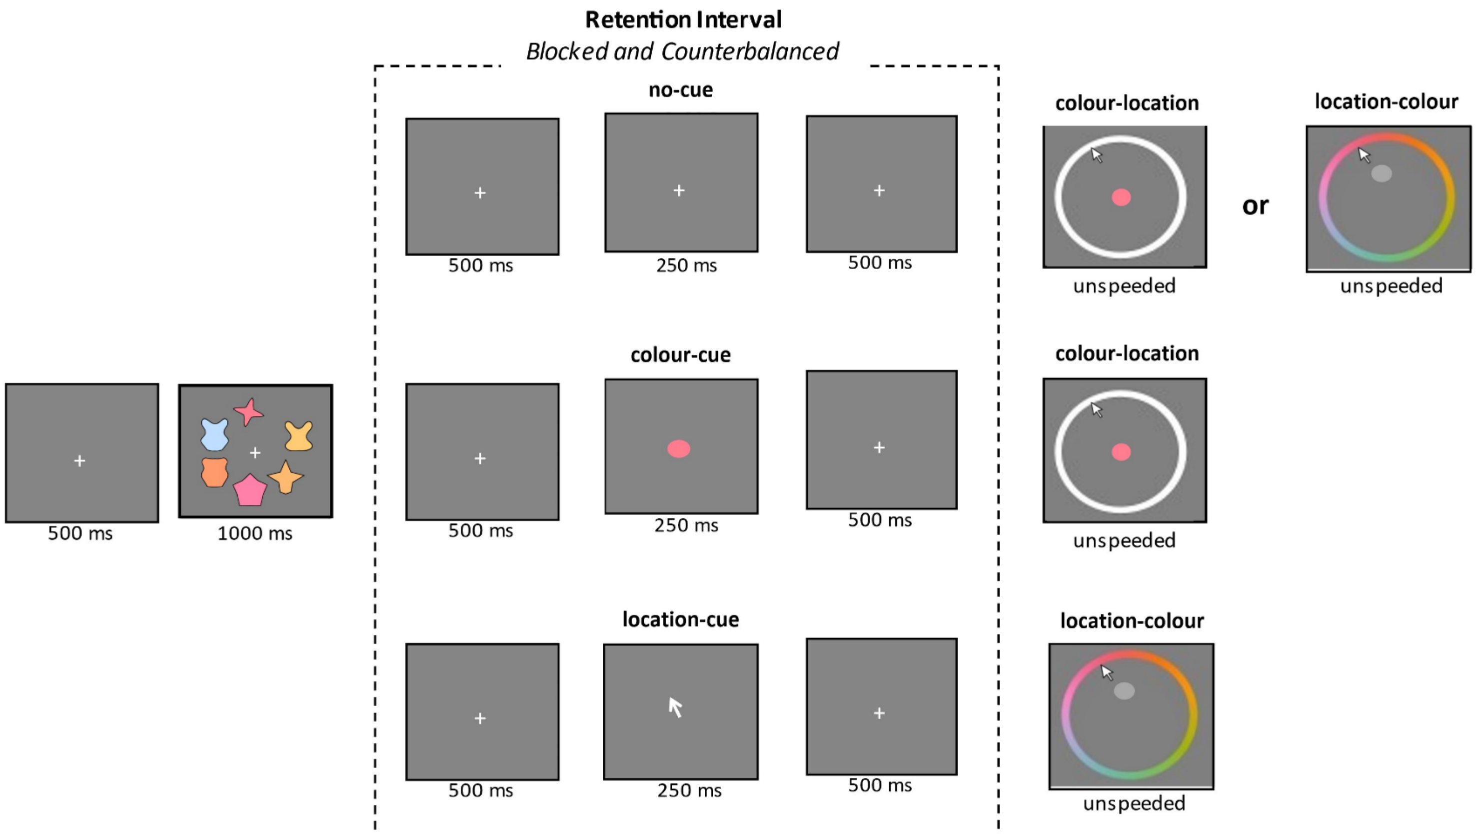

Data Analysis. Analysis scripts to reproduce the results for the experiment are available on the open science framework (OSF; https://tinyurl.com/yc8j5jnn). Observed recall error data collected was fit with a hierarchical Bayesian three-parameter mixture model (Oberauer et al., 2017). The model assumes that the distribution of observed responses reflects the contributions of (i) the probability that the tested feature is held within WM with a (ii) specific precision, and (iii) the probability of misbinding or (iv) guessing when the participant has not stored the information in WM. My hypotheses pertained to binding errors (i.e., the probability of recalling non-target but presented features from the array). The model was fit using rjags (Plummer, 2016) via Markov

Chain Monte Carlo sampling. I verified good convergence and conducted posterior predictive checks to ensure appropriate model fit.

Results

In the current study, the aim was to investigate whether location features are privileged in visual WM, such that guiding attention to location features yields stronger benefits over non-location features during recall from visual WM. Against my first hypothesis, there was no credible difference in binding errors when recalling location compared to colour in the baseline condition (estimated difference = 0.07 [-0.06, 0.19]). This result could be due to the overall low binding errors in the baseline conditions (see Figure 9C). Additionally, I did not observe credible effects in the target memory, precision or guessing parameters (see Table 7 for a full summary). Thus, my first result suggests that there are no differences between retrieving location and colour features. This suggests that while attention is not guided, colour and location features are retrieved similarly and without a bias towards the nature of the to-be-retrieved feature. Therefore it can be argued that location features are not privileged.

Against my second hypothesis, I did not find credible effects in binding errors (see Table 8 for a full summary) irrespective of the nature of the retro-cue (colour or location) or the probe-target type (colour-location or location-colour). This is congruent with the findings of Experiment 2 of Chapter 2. There were, however, credible effects of guiding attention via colour and location retro-cues to target memory and guessing, consistent with Experiment 2 in Chapter 2. Overall, guiding attention via colour and location retro-cues improved target memory and reduced guessing, irrespective of the nature of the retro-cue or the target feature, but had no impact on binding errors.

Congruent with my third hypothesis, I found that binding errors were similar, irrespective of the retro-cue condition (none, colour, and location). This suggests that non-location features are not bound to their respective location features. Further, this suggests that attentional resources are not biased towards location or non-location features.

Table 7

Summary of Mean Differences [and 95% HDIs] in Each Memory Parameter of the Colour-Location vs. Location-Colour Baseline/None Retro-Cue Types.

Probe-target type	Cue-Type	P(Target)	P(Binding Error)	P(Guessing)	Precision
Colour-Location vs. Location-Colour	Baseline/None	0.00 [-0.09, 0.10]	0.07 [-0.06, 0.19]	-0.07 [-0.23, 0.09]	-0.45 [-3.78, 2.79]

Table 8

Summary of Mean Differences [and 95% HDIs] in Each Memory Parameter for Baseline/None vs. Colour and Location Retro-Cue Types.

Probe-target type	Baseline/None vs.	P(Target)	P(Binding Error)	P(Guessing)	Precision
Colour-Location	Colour	**0.19 [-0.09, 0.29]**	0.08 [-0.07, 0.21]	**-0.27 [-0.44, -0.09]**	1.55 [-1.26, 4.24]
Location-Colour	Location	**0.29 [-0.15, 0.44]**	0.04 [-0.13, 0.19]	**-0.34 [-0.55, -0.10]**	3.49 [-0.60, 7.48]

Note: Effects in boldface font indicate credible effects

Figure 9

Posterior Parameter Estimates for the Probability of Target Memory (A), Memory Precision (B), Binding Errors (C), and Probability of Guessing (D). Larger Dark Circles Indicate Group Means, Smaller Faded Circles Indicate Individual Means, and The Error Bars Show the 95% HDIs of the Posterior.

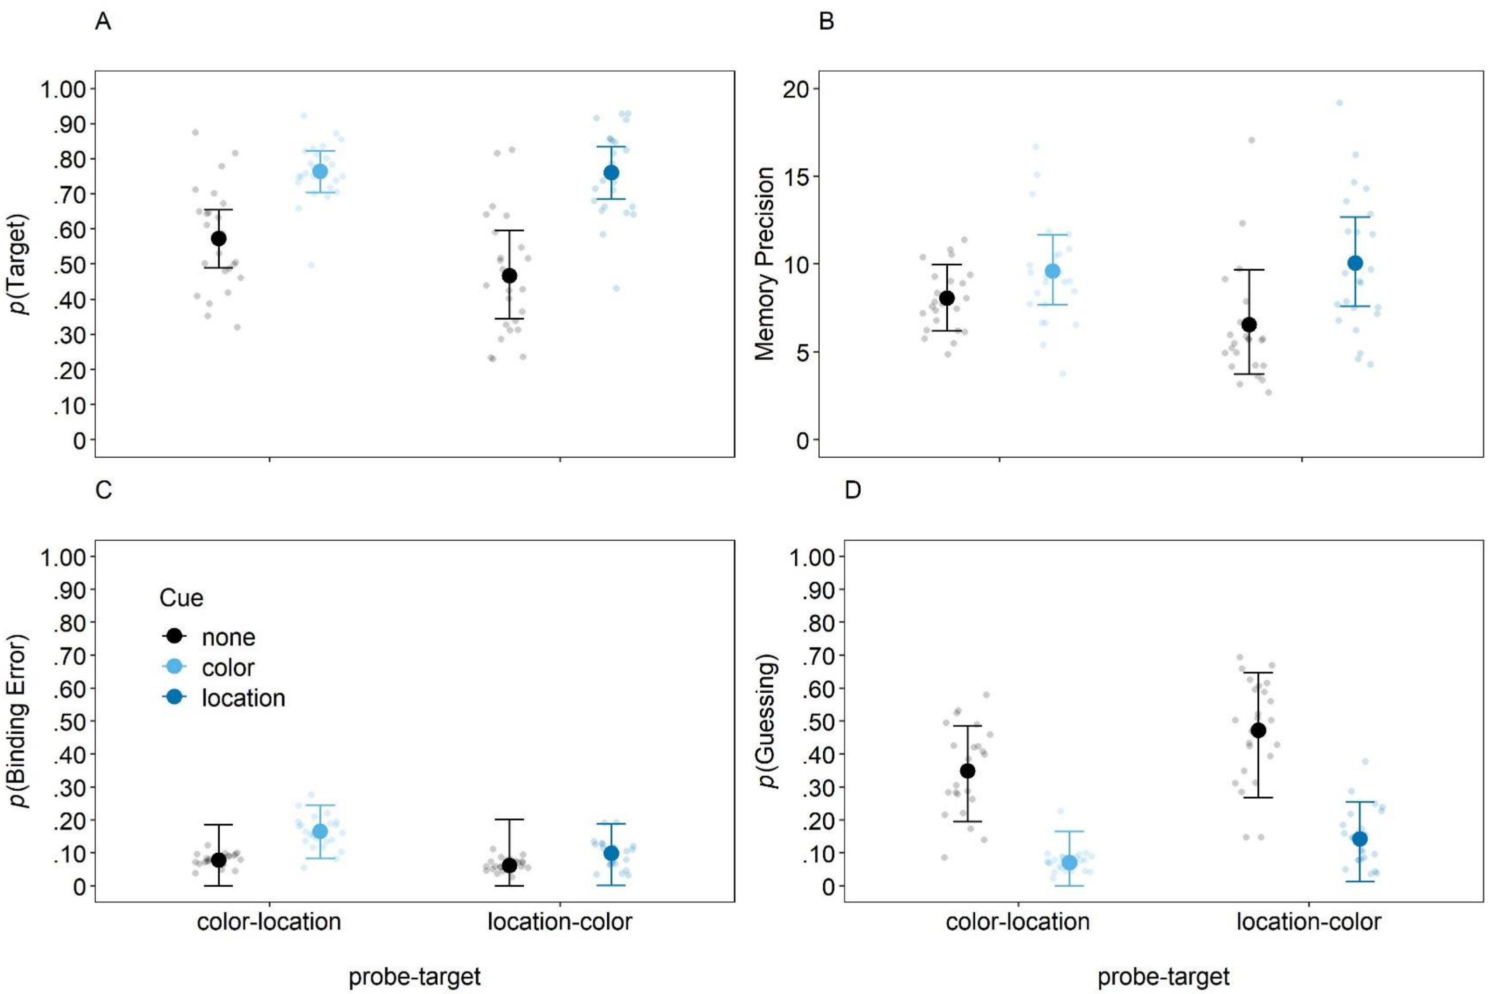

Discussion

Despite prior evidence supporting the notion that location features may be privileged in visual WM (Johnson & Pashler, 1990; Li & Saiki, 2015; Rajsic & Wilson, 2014; Schneegans &

Bays, 2017; Shepherdson et al., 2022), the current findings suggest that there are no observable differences in binding errors between retrieving a location and non-location feature in the no-cue/baseline condition. First, binding errors were not lower when retrieving location versus colour in the no-cue condition, nor were there any differences in target memory, precision, or guessing. These findings therefore may suggest that location features are not privileged compared to colour features. With that being said, the binding errors in the no-cue conditions were low overall, which may have obscured potential differences between colour-location and location-colour binding errors.

In line with this first result, there were no credible retro-cue effects in binding errors for either colour or location, thus, suggesting that guiding attention does not impact binding errors, irrespective of the nature of the retro-cue or the tested feature dimension. However, credible retro-effects in the target memory and guessing parameters were observed regardless of the retro-cue and probe-target type. These findings are aligned with prior work that has shown that retro-cues improve target recall regardless of the nature of the retro-cue (Arnicane & Souza, 2021; Experiment 2 in Chapter 2). Therefore, my findings suggest that guiding attention towards an in-memory feature improves the quality of the in-memory features regardless of whether the to-be-retrieved feature is location or colour. This suggests that visual WM maintenance of location and non-location features benefits from internal attentional resources without a bias towards a single feature dimension. This is inconsistent with FIT (Treisman & Gelade, 1980; Wheeler & Treisman, 2002), as FIT argues that attention is selectively oriented towards location features during visual WM maintenance. If this were the case, while attention was guided via a location retro-cue in the current paradigm should have greatly improved colour target memory compared to location target memory. However, my findings did not suggest that the nature of the retro-cue improved target memory for its associated feature dimension. Therefore it can be argued that during visual WM maintenance, the role of attention is global rather than local, in that, attention is not selective towards individual feature dimensions but rather WM maintenance is benefited from attention by keeping goal-relevant information active. Moreover, it can be argued that while FIT (Treisman & Gelade, 1980) is an influential model in visual perception, and it offers a conceptual framework for the role of attention during feature processing, it may not be the best explanation for feature binding in WM. Further research is required to determine the goal-orientated processing of visual features in WM.

The question remains as to whether visual information is processed as in a feature-based or object-based manner. In line with my final hypothesis, my findings suggest that the lack of observable differences in binding errors between location and non-location features may indicate that visual information may be processed as bound representations (thus agreeing with the object-based view). This suggests that multi-feature objects are represented in a bound state in visual WM. Arguably, this suggests that guiding attention via retro-cues strengthens the maintenance of in-memory bindings, thereby improving target memory. The current findings are also aligned with prior research that has demonstrated that attentional resources can be selectively guided towards a target feature to aid the maintenance of visual features within WM and improve recall performance (Hajonides et al., 2020; Astle et al., 2012; Li & Saiki, 2015; see Souza & Oberauer, 2016 for a review). In the current work, it was observed that retro-cues improved target memory and reduced guessing, irrespective of the nature of the target. It can therefore be argued that while the processing of visual information may be object-based and without a definitive bias towards location-based features, selectively guiding internal attentional resources towards an in-memory features strengthens their maintenance.

In summary, I investigated whether the processing of location features is privileged in visual WM over non-location features by manipulating the probe-target combination during retrieval and by guiding attention via feature-specific retro-cues. Against the findings of Rajsic and Wilson's (2014; Experiment 1), I did not find any evidence to suggest that there are differences in binding errors between non-location (colour) and location features in visual WM. Moreover, I did not observe that guiding attention via location and colour retro-cues reduced colour and location binding errors, respectively, I did observe a benefit to target memory and a reduction in guessing while attentional resources were guided via both colour and location retro-cues. Taken together, my findings suggest that visual information may be processed in an object-based manner, with no differences between the to-be-retrieved feature type, but the representations can be strengthened by guiding attention via retro-cues.

Chapter 5

Temporary removal underlies local working memory updating.

Rationale

Working memory (WM) updating is a core ability that allows one to replace outdated (irrelevant) information with new (relevant) information to be kept active during goal-orientated processing (Ecker, Oberauer, et al., 2014). For example, when shopping for a white pair of trainers, you may notice a red pair that are similar in style to the white pair and fixate on them to plan a new outfit. This simple example demonstrates how we can update our WM in a real-life situation, wherein outdated information is removed to prioritize updated information (see Lewis-Peacock et al., 2018). This raises the question of whether individual elements of information are updated in WM (e.g., just the colour of the sneaker) or whether the entire contents of WM are updated (the whole representation of the sneaker). Moreover, how may outdated information continue to influence WM after it has been updated?

Prior research has suggested two parameters involved in WM updating: local updating and global updating (Kessler & Meiran, 2008). Local updating occurs when specific items can be independently maintained and updated irrespective of other information in WM. Conversely, global updating occurs when all the retained items are updated even if only one of them needs modifying (Kessler et al., 2015; Kessler & Meiran, 2008). Kessler and colleagues (2015) investigated whether local or global updating processes occur during partially or fully updating the contents of visual WM during a change detection task. During the task, participants were presented with a set of coloured squares, followed by one of several intervals: a fixation cross (baseline); the same set of coloured squares for a second time (repeated); a set of coloured squares where one or more colours had updated (partial-update, i.e., 1/6, 2/6, or 4/6 squares updated in colour); or a new set of coloured squares (full-update, i.e., 6/6 squares updated in colour). The participants then determined whether a probe was the same or different in colour as in the last-presented array, which could have been the location of a repeated-item or an updated-item depending on the update condition. The findings suggested that visual WM updating may be local given the greater recall accuracy during the repeated-location probes compared to the updated-location probes, regardless of the number of items to update. If global updating had occurred instead, then all the items presented in the partial-update and full-update conditions should have been encoded as new, in which case there would be no observable recall benefit to probing repeated over updated locations. Thus, visual WM updating may predominantly rely on local rather than global updating, such that information is independently updated regardless of other information in WM.

Although Kessler and colleagues' research highlighted that visual WM updating may rely on local updating, it is unclear as to what happens to outdated information during updating. In a recent review, Lewis-Peacock and colleagues (2018) proposed that the removal of outdated information from WM could be either temporary or permanent. The authors posited that temporary removal momentarily withdraws outdated information, allowing goal-relevant information to be prioritized, but the outdated information may still be accessible if required for future processing. Conversely, permanent removal of outdated information relies on its withdrawal from WM when it is no longer goal relevant. The authors postulated that the removal of information from WM can occur through the removal of its content, context, or the content-context bindings. Supporting Lewis-Peacock and colleagues' removal hypotheses, prior research has suggested that the removal of outdated information may occur through the unbinding of item-position associations (Ecker, Lewandowsky, et al., 2014). Further research has demonstrated that temporarily removing outdated information (by means of suppression) during WM updating hinders the prioritization of goal-relevant information (Palladino et al., 2001) and reduces its recall accuracy (Carriedo et al., 2016). Additionally, neuroscientific research has shown that the neural signatures for task-irrelevant information decrease to baseline while outdated information is no longer prioritized in the focus of attention, but are decodable once again after the temporarily removed information returns to the focus of attention (Lewis-Peacock et al., 2012; Lewis-Peacock & Postle, 2012). Furthermore, neural patterns for outdated information decrease during its active removal from WM (Kim et al., 2020). While this evidence demonstrates that outdated information may be temporarily or permanently removed, prior work did not investigate which form of removal occurs during local WM updating. Therefore, given Kessler and colleagues' (2015) suggestion that WM updating is local, it is of interest to understand whether outdated information is temporarily or permanently removed during local WM updating.

Current Study

I adapted Kessler and colleagues' (2015) paradigm to address whether local updating requires the temporary or permanent removal of outdated information by using a continuous-report measure during the retrieval phase. Using a continuous-report paradigm builds on the prior work of Kessler and colleagues by revealing how removal impacts the mixture of components underlying observed recall error (i.e., the distance between the target colour and the participant's response). Accordingly, we fit two hierarchical Bayesian three-parameter mixture models to recall error to estimate parameters of target recall, precision, binding errors, and guessing (e.g., Loaiza & Souza, 2019; Oberauer et al., 2017). Model 1 simulated the temporary removal of outdated information (via suppression) by including the irrelevant (outdated) and relevant (new) features during model fitting. Model 2 simulated the permanent removal of irrelevant information by including only the relevant (new) features during model fitting. I applied a widely applicable information criterion (WAIC; Oberauer et al., 2017; Watanabe, 2021) to compare which model fit the data best.

I hypothesized that if local WM updating requires the temporary removal of irrelevant information, then Model 1 will best fit the partial-updating conditions, and binding errors should occur most often during the partial-update conditions compared to the baseline, repeated and full-update conditions. This should be particularly evident when probed with an updated item versus a repeated item. Conversely, if local WM updating relies on permanent removal, then Model 2 will best fit the full-update condition and binding errors should be similar across the partial- and full-update conditions. These hypotheses pertain to binding errors, specifically given Lewis-Peacock and colleagues' (2018) suggestion that temporary removal may involve unbinding item-position associations, in which case the irrelevant content may still be accessible and interfere with recall.

Method

Participants. In line with Kessler and colleagues' (2015) paradigm, 30 complete and valid datasets were collected from participants aged 18-35 through Prolific who were compensated with UK £7.50/hour. All participants reported normal or corrected-to-normal vision and colour vision, and the latter was verified with a colour blindness task. An additional 18 participants were excluded. 15 excluded participants did not meet the minimum colour vision requirements. One excluded participant did not continue with the current study after the completion of the colour blindness task. Two excluded participants did not complete the current study.

Materials and Procedure. The experiment was conducted online through lab.js (Henninger et al., 2022) and was hosted on the JATOS server, Mindprobe (https://jatos.mindprobe.eu; Lange et al., 2015). The experiment was conducted in two phases: Participants who successfully completed Phase 1 (demographic questionnaire and colour blindness test) were redirected to Phase 2 (main task), whereas those who did not meet the inclusion criteria were debriefed and compensated for their time up to that point. During Phase 2, the 400 total trials of the visual WM updating task followed a similar format (see Figure 10): After a fixation cross was presented for 600 ms, six coloured dots (of radius = 25°, arranged equidistantly around an invisible circle with a radius of 150°) were presented for 100 ms and followed by a 600 ms retention interval. The colours were sampled without replacement from 360 possible CIE*L*a*b colours (L = 70, a= 20, b = 38). Thereafter, in half (200) of the trials, participants were immediately directed to the retrieval phase (baseline condition), whereas the remaining trials presented a second array of coloured dots for 100 ms (update conditions: repeated, partial-update, or full-update). The trials of the baseline/update conditions were randomly intermixed. In the repeated condition, participants were presented with the original array of coloured dots for a second time. In the partial-update conditions (1/6, 2/6, 4/6), some of the colours from the first array were updated in the second array (e.g., in the 1/6 partial-update condition, only one of the colours was updated and the remaining five were repeated). During the full-update (6/6) condition, all six colours of the second array were updated.

During retrieval, participants were probed with a grey dot in the location of a prior presented dot and were required to recreate the colour of the probed location by moving their mouse along a colour wheel (radius = 250°). The probe was, by necessity, a repeated item in the baseline and repeated conditions, an updated item in the full-update condition, and a repeated or updated item with equal probability in the partial-update conditions. In the repeated, partial- and full-update conditions, participants were instructed to respond to the most recently presented array of coloured dots (i.e., the repeated or updated dots). Participants completed 10 blocks of trials, with a break offered after 40 trials. After every 10 trials, participants received feedback about their average recall accuracy (expressed as a percentage of their mean reproduction error, i.e., 100 – 100 * mean error/180). Before completing the 400 test trials, participants completed 6 practice trials and were provided with feedback about their average recall accuracy.

Figure 10

Example of the Task Procedure. From Top to Bottom, Baseline Condition, Update 0/6 (Repeated) Condition, Update 1/6, 2/6, or 4/6 Condition (Repeated-Item Location Probe, Updated-Item Location Probe), Update 6/6 Condition (Updated-Item Location Probe).

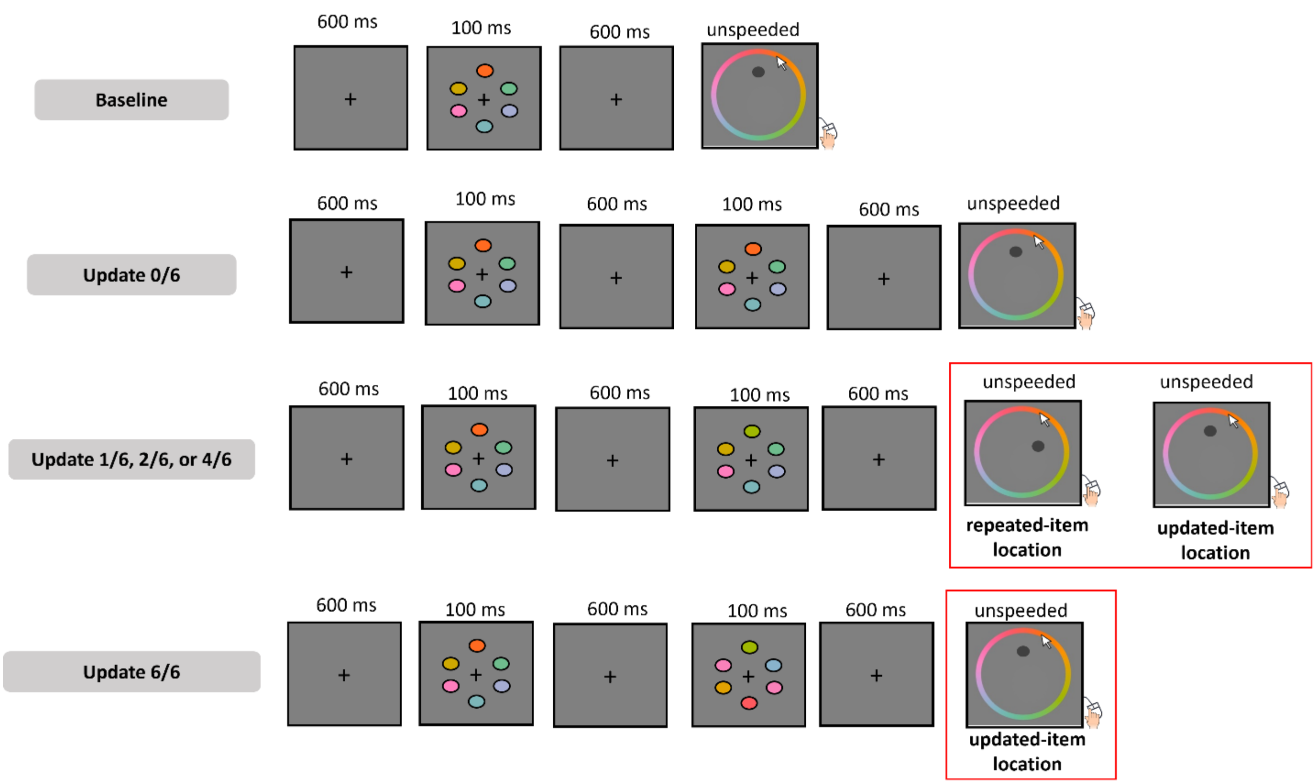

Design and Data Analysis. The main independent variable of update conditions (baseline, 0/6, 1/6, 2/6, 4/6, and 6/6 updates) was manipulated within-subjects, and the probe location (repeated or update) necessarily nested within each update condition (baseline and 0/6: repeated only; partial updates: repeated and updated; full 6/6 update: updated only). Although the hypotheses did not pertain to the task performance between the update conditions, we followed Kessler and colleagues' (2015, Experiment 1) analyses to determine whether the update conditions and probe locations impacted recall error by conducting a Bayesian analysis of variance (BANOVA) using the R package BayesFactor (Morey et al., 2022). Using a Bayesian approach allows for a comparison between two models based on the Bayes Factor (BF) statistic, which conveys the ratio between the likelihoods of the data under the two compared models. For example, to determine the evidence for a main effect of update condition, one would compare the likelihood of a model assuming an effect (M_1) relative to the null model (M_0; BF_{10}). BFs are interpreted continuously, such that, a BF between 1 to 3 (or 0.1 to 0.3) is interpreted as ambiguous, and BFs greater than 3 and 10 (0.30 and 0.10) are substantial and strong evidence for/against the model in the numerator relative to the model in the denominator.

To examine the hypotheses, observed recall error was fit with two different versions of a hierarchical Bayesian three-parameter mixture model (Oberauer et al., 2017). This model assumes that the distribution of observed responses reflects a mixture of different contributions: (i) the probability that the tested feature is held within working memory, (ii) with a specific precision, and the probability of (iii) binding errors or (iv) guessing. Model 1 simulated temporary removal (via suppression) of outdated feature information by including outdated and new, updated, colour information during model fitting. Model 2 simulated permanent removal by only including the updated colour information during model fitting. The two models were fit using rjags (Plummer et al., 2022) via Markov Chain Monte Carlo sampling. I verified good convergence by visually inspecting the plotted chains, and I conducted posterior predictive checks to ensure appropriate model fit by verifying that the R-hat ($\hat{R}$) statistic in each model fit was lower than 1.06. The $\hat{R}$

statistic in both models fit were lower than 1.06, thus indicating that the models fit could adequately discriminate the data that was fit to each model. Posterior estimates of each parameter are reported in Figures 2 and 3, and the posterior differences between the conditions for each parameter (and their 95% HDIs) in Table 2. The anonymized raw data and analysis scripts are openly accessible on the Open Science Framework (OSF; https://tinyurl.com/yuzey69).

Results

Following Kessler and colleagues' (2015) analyses, I conducted two BANOVAs on recall error (see Table 9) given that the independent variables were not fully crossed. First, I investigated whether the presentation of a second memory array affected recall error in a one-way within-subjects BANOVA comparing the baseline, 0/6, and 6/6 update conditions. Consistent with Kessler and colleagues, I observed strong evidence that the update condition impacted recall error, $BF_{Update/Null} = 105.96$. Follow-up tests showed ambiguous evidence against a difference between the baseline and 6/6 conditions, $BF_{10} = 0.53$. However, inconsistent with Kessler and colleagues, repeating the array twice in the 0/6 condition did not substantially impact recall error compared to the baseline ($BF_{10} = 0.52$) and 6/6 ($BF_{10} = 2.12$) conditions. Next, a 3 (partial-updates: 1/6, 2/6, 4/6) x 2 (probe-location: repeated, updated) within-subjects BANOVA on recall error showed evidence against an effect of partial-update conditions, BF = 0.21, and the interaction, BF = 0.20, similar to Kessler and colleagues' findings. However, incongruent with their results, I did not observe that repeated probe locations yielded lower recall error than updated locations, BF =0.12. Overall, the evidence for any effects of update condition or probe location on recall error were weak compared to Kessler and colleagues' original findings, and thus the hypothesized modelling of the processes underlying local WM will be helpful.

Table 9

Means (and Standard Deviations) of Recall Error as a Function of Update Condition.

	Update condition					
Probe location	Baseline	0/6	1/6	2/6	4/6	6/6
Repeated	67.31 (12.09)	62.42 (16.99)	66.72 (22.96)	70.37 (18.61)	67.85 (19.12)	-
Updated	-	-	66.77 (16.20)	70.88 (17.07)	70.77 (17.69)	72.08 (16.18)

To examine the hypotheses, I considered two models in our analysis: Model 1 simulated the temporary removal of irrelevant colour features by including the irrelevant (outdated) colours and the relevant (new) colours during model fitting, whereas Model 2 simulated the permanent removal of irrelevant colours by only including the relevant (new) colours during model fitting. I compared the model fits through applying a widely applicable information criterion (WAIC; Gelman et al., 2014; Oberauer et al., 2017; Watanabe, 2021), which indicated that Model 1 best fit both the partial-updating and full-update conditions (see Table 10). These results support the hypothesis that temporary removal of outdated (irrelevant) best fits the observed responses during the partial-update conditions, as well as the full-update condition. Thus, local WM updating may rely on the temporary removal of irrelevant information regardless of the number of updates.

Table 10

WAIC Values for Each Updating Condition (Partial and Full), Per Model Simulating Temporary and Permanent Removal.

Model	Simulating	Updating Condition	WAIC
1	Temporary Removal	1/6 partial update	2.14
		2/6 partial update	1.42
		4/6 partial update	22697.74
		6/6 full update	-9195.96
2	Permanent Removal	1/6 partial update	33674.80
		2/6 partial update	31061.11
		4/6 partial update	36826.78
		6/6 full update	-4944.28

Note. Lower WAIC values indicate better model fit.

The posterior parameter estimates of Model 1 are presented in Figure 11, and the differences of the partial-update conditions relative to the baseline, repeated (0/6) and full-update (6/6) conditions are presented in Table 11. The results showed credibly greater binding errors in the 2/6 and 4/6 partial-update conditions relative to the baseline condition for both repeat and update probe locations. There were also credibly greater binding errors in the full-update (6/6) condition relative to the 1/6 partial-update condition as well as compared to the baseline (latter not presented in Table 4; estimate = 0.29 [0.13, 0.44]). This suggests that binding errors became more prevalent with an increased number of items to update. There were also a couple of credible differences in the probabilities of guessing and recalling the target for repeat probes that were not as systematic, and so I refrain from any strong interpretation. The remaining comparisons were not credible.

Although Model 1 fit the data best, for the sake of comparison and completeness, I also report on the parameter estimates and differences of Model 2 in Figure 12 and Table 12. The striking difference is that there were no longer any credible differences in binding errors between the baseline and partial-update conditions in Model 2 as was the case in Model 1. This reinforces the notion that the better fit of Model 1 that simulates temporary removal likely results from the residual traces of the updated memory items that lead to binding errors, which is not captured in Model 2 that simulates permanent removal.

Table 11

Model 1 Summary of Mean Differences [and 95% HDIs] in Each Memory Parameter of Partial-Update Conditions versus the Baseline, Repeated and Full-Update Conditions.

Comparison	Partial-Update Condition	Probe Location	*P*(Target)	*P*(Binding Error)	*P*(Guessing)	Precision
Baseline vs.	1/6	Repeat	0.03 [-0.10, 0.16]	0.08 [-0.07, 0.23]	-0.11 [-0.28, 0.06]	1.12 [-2.46, 4.88]
		Update	0.04 [-0.08, 0.15]	0.13 [-0.01, 0.27]	-0.17 [-0.35, 0.01]	2.54 [-2.34, 8.01]
	2/6	Repeat	-0.01 [-0.12, 0.11]	**0.21 [0.06, 0.35]**	**-0.20 [-0.37, -0.04]**	-0.32 [-4.13, 3.85]
		Update	-0.03 [-0.14, 0.07]	**0.18 [0.03, 0.34]**	-0.15 [-0.33, 0.04]	1.91 [-3.11, 7.51]
	4/6	Repeat	0.01 [-0.11, 0.11]	**0.21 [0.07, 0.36]**	**-0.22 [-0.40, -0.06**]	3.27 [-2.31, 9.45]
		Update	-0.06 [-0.17, 0.05]	**0.18 [0.00, 0.37]**	-0.13 [-0.36, 0.09]	0.52 [-3.72, 5.19]
Repeated (0/6) vs.	1/6	Repeat	-0.03 [-0.16, 0.12]	-0.01 [-0.16, 0.17]	0.03 [-0.17, 0.20]	1.58 [-2.34, 5.38]
		Update	-0.03 [0.15, 0.10]	0.06 [-0.15, 0.10]	-0.03 [-0.22, 0.18]	2.99 [-2.25, 8.42]
	2/6	Repeat	-0.07 [-0.19, 0.06]	0.13 [-0.03, 0.30]	-0.07 [-0.25, 0.12]	0.14 [-3.72, 4.54]
		Update	-0.09 [-0.21, 0.04]	0.11 [-0.06, 0.28]	-0.01 [-0.20, 0.20]	2.86 [-2.23, 8.72]
	4/6	Repeat	-0.05 [-0.17, 0.06]	0.14 [-0.03, 0.30]	-0.08 [-0.26, 0.10]	3.73 [-1.95, 10.37]
		Update	-0.12 [-0.23, -0.01]	0.11 [-0.09, 0.31]	0.01 [-0.23, 0.24]	0.98 [-3.63, 5.93]
Full Update (6/6) vs.	1/6	Repeat	0.08 [-0.06, 0.22]	**-0.21 [-0.39, -0.01]**	0.13 [-0.10, 0.33]	0.85 [-3.97, 5.31]
		Update	0.09 [-0.04, 0.21]	-0.16 [-0.35, 0.03]	0.07 [-0.18, 0.29]	2.29 [-3.43, 8.37]
	2/6	Repeat	0.05 [-0.07, 0.17]	-0.08 [-0.28, 0.11]	0.03 [-0.19, 0.24]	-0.59 [-5.03, 4.62]
		Update	0.02 [-0.10, 0.14]	-0.11 [0.30, 0.08]	0.09 [-0.14, 0.31]	2.13 [-3.57, 8.50]
	4/6	Repeat	0.06 [-0.06, -0.17]	-0.08 [-0.27, 0.11]	0.02 [-0.19, 0.23]	3.00 [-3.28, 9.69]
		Update	0.00 [-0.12, 0.12]	-0.11 [-0.34, 0.10]	0.11 [-0.15, 0.36]	0.25 [-4.80, 5.74]

Note. Effects in boldface font indicate credible effects.

Figure 11

Posterior Parameter Estimates of Model 1 for the Probability of Target Memory (A), Memory Precision (B), Binding Errors (C) and Probability of Guessing (D). Larger Dark Circles Indicate Group Means, Smaller Faded Circles Indicate Individual Means, and the Error Bars Show the 95% HDIs of the Posterior.

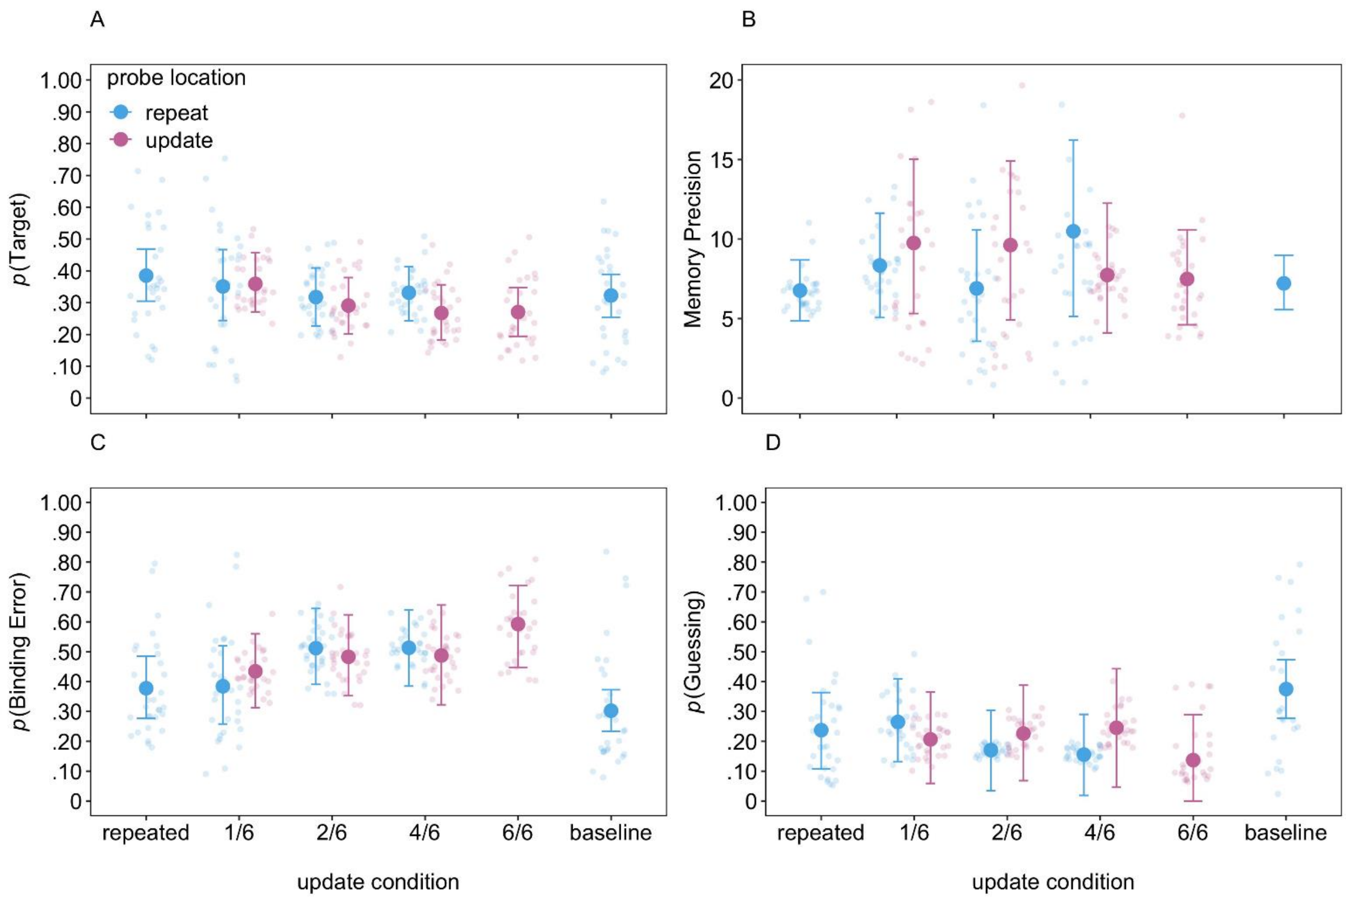

Table 12

Model 2 Summary of Mean Differences [and 95% HDIs] in Each Memory Parameter of Partial-Update Conditions versus the Full-Update (6/6) Condition.

Comparison	Partial-Update Condition	Probe-Location	*P*(Target)	*P*(Binding Error)	*P*(Guessing)	Precision
Baseline vs.	1/6	Repeated	0.03 [-0.10, 0.16]	0.04 [-0.11, 0.19]	-0.07 [-0.25, 0.10]	0.91 [-3.06, 5.12]
		Update	0.04 [-0.07, 0.15]	0.06 [-0.08, 0.20]	-0.10 [-0.28, 0.07]	3.14 [-2.19, 9.00]
	2/6	Repeated	0.00 [-0.12, 0.12]	0.08 [-0.05, 0.21]	-0.08 [-0.24, 0.10]	0.72 [-3.68, 6.02]
		Update	-0.03 [-0.15, 0.09]	0.08 [-0.07, 0.24]	-0.05 [-0.25, 0.16]	2.47 [-3.12, 8.83]
	4/6	Repeated	-0.02 [-0.14, 0.11]	-0.02 [-0.16, 0.12]	0.03 [-0.17, 0.24]	4.65 [-1.98, 12.53]
		Update	-0.04 [-0.16, 0.07]	0.06 [-0.09, 0.22]	-0.02 [-0.23, 0.19]	-0.38 [-4.61, 4.36]
Repeated (0/6) vs.	1/6	Repeated	-0.03 [-0.17, 0.11]	-0.04 [-0.19, 0.14]	0.06 [-0.12, 0.26]	1.40 [-2.65, 5.85]
		Update	-0.02 [-0.15, 0.10]	-0.01 [-0.17, 0.14]	0.04 [-0.16, 0.24]	3.63 [-2.02, 9.56]
	2/6	Repeated	-0.06 [-0.19, 0.06]	0.00 [-0.14, 0.15]	0.06 [-0.11, 0.25]	1.21 [-3.65, 6.25]
		Update	-0.09 [-0.22, 0.03]	0.01 [-0.17, 0.17]	0.09 [-0.13, 0.29]	2.96 [-2.72, 9.30]
	4/6	Repeated	-0.08 [0.21, 0.06]	-0.09 [-0.25, 0.07]	0.17 [-0.05, 0.40]	5.14 [-1.97, 12.69]
		Update	-0.10 [-0.24, 0.02]	-0.01 [-0.18, 0.16]	0.11 [-0.12, 0.33]	0.11 [-4.10, 5.24]
Full Update (6/6) vs.	1/6	Repeated	0.10 [-0.04, 0.22]	0.03 [-0.15, 0.20]	-0.13 [-0.32, 0.08]	1.33 [-3.22, 6.18]
		Update	0.10 [-0.02, 0.22]	0.06 [-0.11, 0.21]	-0.16 [-0.37, 0.06]	3.56 [-2.41, 9.83]
	2/6	Repeated	0.06 [-0.06, 0.18]	0.07 [-0.08, 0.22]	-0.13 [-0.32, 0.05]	1.15 [-4.15, 6.70]
		Update	0.03 [-0.09, 0.15]	0.08 [-0.11, 0.25]	-0.11 [-0.33, 0.13]	2.89 [-3.33, 9.76]
	4/6	Repeated	0.05 [-0.08, 0.18]	-0.02 [-0.19, 0.14]	-0.02 [-0.25, 0.21]	5.07 [-1.86, 13.28]
		Update	0.02 [-0.09, 0.14]	0.06 [-0.10, 0.23]	-0.08 [-0.31, 0.13]	0.04 [-4.71, 5.37]

Figure 12

Posterior Parameter Estimates of Model 2 for the Probability of Target Memory (A), Memory Precision (B), Binding Errors (C) and Probability of Guessing (D). Larger Dark Circles Indicate Group Means, Smaller Faded Circles Indicate Individual Means, and the Error Bars Show the 95% HDIs of the Posterior.

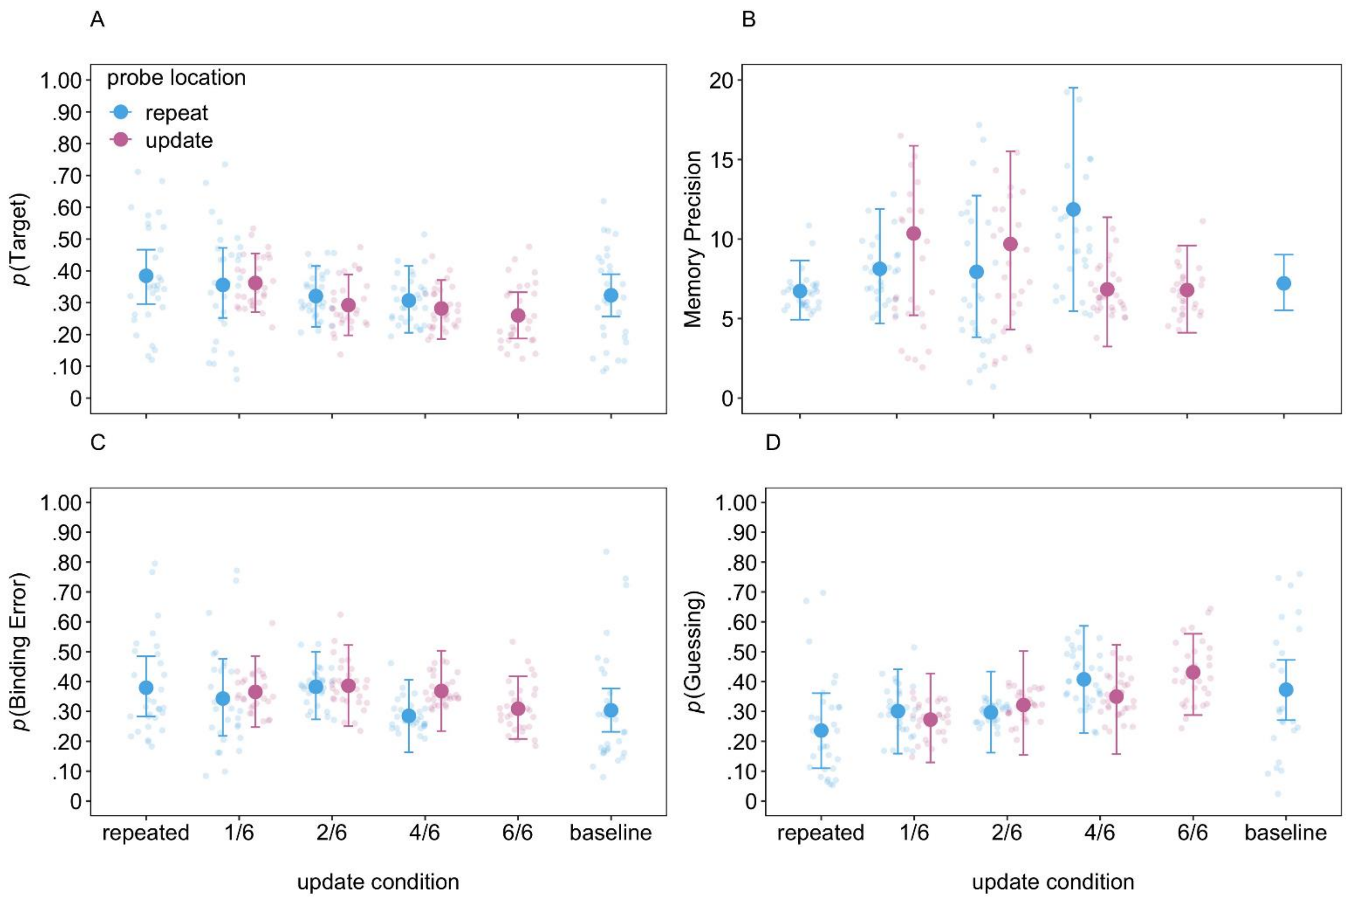

Discussion

In the current study, I investigated whether local WM updating relies on the temporary or permanent removal of outdated (irrelevant) information. To reiterate, Lewis-Peacock and colleagues (2018) proposed that during temporary removal, outdated information is momentarily withdrawn from WM, but it is still accessible if the outdated information is required for future goal-orientated processing. On the other hand, during permanent removal, outdated information is no longer active within WM, and it cannot be accessed for future processing. I investigated whether temporary or permanent removal underlies local WM updating by adapting a visual WM updating task (Kessler et al., 2015) and comparing the fit of the data to two versions of a three-parameter mixture model (Bays et al., 2009; Oberauer et al., 2017) that simulated temporary or permanent removal. The findings suggested that the model simulating temporary removal of outdated information best explained local WM updating, as evidenced by lower WAIC values in both the partial- and full-updating conditions compared to the WAIC values of the model simulating permanent removal. The parameter estimates of Model 1 further reinforced the notion that outdated information still impacted observed recall error given that binding errors were credibly higher during the partial-update conditions compared to the baseline, whereas this was not the case for Model 2.

These findings suggest that opportunities for binding errors increase when outdated information is still accessible in WM. Specifically, binding errors were more frequent during the 2/6 and 4/6 partial-update conditions, as well as the 6/6 full-update condition, relative to the baseline. This suggests that there are greater opportunities for new feature information to become misbound as the outdated information remaining in WM increases. This pattern was not observable with recall error alone. Indeed, contrary to Kessler and colleagues' (2015) findings, I did not observe that repeating an array (0/6 vs. baseline and 6/6) nor repeated probes (vs. updated probes) impacted task performance. One explanation of this may be that participants failed to update WM during the current study: If this were the case, recall error (Table 9) would have been much higher than I have observed. This highlights that a modelling approach allows for a closer examination of underlying WM updating processes underlying task performance. That is, although I did not observe evidence for local updating in recall performance, applying a mixture model revealed that temporary removal most strongly impacts binding errors. Overall, the current findings indicate that visual information can be locally updated in WM, however this comes with a cost when outdated information is temporarily removed, as evidenced by the increase in binding errors while more than one item is to be updated. Therefore, this highlights that participants were indeed able to update the contents of their WM, but more powerful statistical measures are required to understand the underlying processes (such as the models fit in the current study).

Furthermore, I did not find credible differences in binding errors while comparing the partial-update conditions to the repeated condition, and I only observed a credible difference in binding errors in the 1/6 partial-update condition relative to the full-update condition. Thus, the differences between and within the partial-update and full-update conditions were less evident than when comparing those conditions to a baseline, where no new information was presented. These findings suggest that the presence of outdated information may interfere with the binding of new (updated) information to its associated feature, thus causing more opportunities for binding errors to occur and irrespective of whether the binding is partially or fully updated. Conversely, I found that binding errors were similar in Model 2 (simulating permanent removal), irrespective of the number of items that were to be updated. This suggests that the pattern of binding errors of Model 1 was specific to the outdated items, which yielded a better fit to the data compared to Model 2.

Thus, these results collectively suggest that local WM updating relies on the temporary removal of outdated information.

One could argue that during temporary removal, while outdated information is still accessible in WM, the prioritization of new (updated) bindings is fickle and outdated information may interfere with one's ability to bind new information together. This concurs with prior research that has argued that irrelevant information interferes with the binding of new information during WM updating. For example, Szmalec and colleagues (2011) argued that during WM updating, to-be-retained bindings may be influenced by proactive interference from irrelevant information (i.e., the phenomena in which prior information interferes with new learning). The current findings from the model simulating temporary removal are consistent with these accounts, such that while outdated information is temporarily removed (but still accessible), binding new information to in-memory information comes at a cost: The presence of outdated information harms the maintenance of correct updated-retained bindings (e.g., a new colour is bound to its associated location). In the case of temporary removal, not only are updated-retained bindings present in WM, but arguably outdated-retained bindings (e.g., an outdated colour and its associated location) are also accessible (although outdated-retained bindings may be suppressed or inhibited). Therefore, prioritizing the new-retained binding while the outdated-retained binding is still accessible leads to source confusion during recall, in turn, the new-retained binding is confused with the outdated-retained binding during prioritization and thus, opportunities for binding errors occur. Future research may consider applying Oberauer and colleagues' (2017) Bayesian hierarchical interference mixture model to understand whether interference of outdated-retained bindings truly impacts new-retained bindings.

On the contrary, in the model simulating permanent removal, we observed that binding errors were similar irrespective of the update condition and the probe location. Lewis-Peacock and colleagues (2018) proposed that during permanent removal, outdated information is unbound from its context (i.e., content-context binding). The findings of the model simulating permanent removal suggests that this may be the case: While outdated-retained content-context bindings are not present in WM, then only the updated-retained content-context bindings are prioritized. Prior work has also argued that outdated information is removed through the unbinding item-position bindings (Ecker, Lewandowsky, et al., 2014), and location-based features (such as an item's position) are contextual (Ecker et al., 2013). Thus, it can be argued that during permanent removal, prioritizing updated-retained bindings relies on the prioritization of contextual information. This suggests that during permanent removal, retained contextual information is favoured over the updated content. Through this prioritization, updated information may be directly bound to its respective context and without outdated-retained bindings hindering the direct binding. Thus, the direct binding between updated and contextual information may explain why fewer binding errors occur during permanent removal.

In summary, the current chapter demonstrates that the temporary removal of outdated information may best explain local WM updating. My findings show that while outdated information is still accessible in WM, binding errors increase with the increased number of items that are updated. I interpret this to suggest that the binding of updated and retained information is fickle and outdated-retained bindings may interfere with this. Conversely, the model simulating permanent removal indicated that binding errors are similar, irrespective of the number of items that are to be updated. This may be explained by the prioritization of an item's contextual features during permanent removal and one's ability to bind updated information to its respective context.

Chapter 6

Locally updating shapes is more difficult than locally updating colours.

Rationale

Following the findings of Chapter 5, I examined whether updating an alternative feature dimension (i.e., the shape of an item) would produce the same pattern of results as that of Chapter 5. If shape features are processed in the same manner as colour features (Chapter 5), then the findings of the previous chapter should be replicated when using recalling shapes. This finding would be aligned with my findings from the previous chapters (Chapter 2-4), wherein I observed that the processing and maintenance of visual information in WM is similar, irrespective of the features' nature. Moreover, this would indicate that temporary removal of outdated information occurs during local WM updating, irrespective of the nature of the to-be-updated feature dimension. The hypotheses remained the same as that of Chapter 5: If local WM updating requires the temporary removal of irrelevant information, then Model 1 will best fit the partial-updating conditions, and binding errors should occur most often during the partial-update conditions compared to the baseline, repeated and full-update conditions. This should be particularly evident when probed with an updated item versus a repeated item. Conversely, if local WM updating relies on permanent removal, then Model 2 will best fit the full-update condition. If local WM updating relies on the permanent removal of irrelevant features, then binding errors should be similar across the partial- and full-update conditions.

Method

Participants. In line with Chapter 5, 30 complete and valid datasets were collected from participants aged 18-35 through Prolific who were compensated with UK £7.50/hour. All participants reported normal or corrected-to-normal vision and colour vision, and the latter was verified with a colour blindness task before starting the experiment. Seventeen participants were excluded from the current study due to not meeting the minimum colour vision requirement.

Materials, Procedure, and Analysis. The materials, procedure and analysis were identical to Chapter 5 with the following exceptions: The current study manipulated whether shapes presented in grey were required to-be-updated in the update conditions rather than coloured dots (see Chapter 5). The six shapes were randomly sampled from 360 shapes (Li et al., 2020) with a minimum distance of 60° between them to reduce the similarity between them. During the retrieval phase, a shape wheel was presented rather than a colour wheel. Depending on the update condition, participants were probed on a repeated-item location or an updated-item location and were required to retrieve the corresponding shape by moving their mouse along the shape wheel. The anonymised raw data, analysis scripts and experiment materials are available on the open science framework (OSF; https://tinyurl.com/yc275d9s).

Results

In line with Chapter 5 and Kessler and colleagues (Experiment 1, 2015), I conducted a two BANOVAs on recall error (Table 13). First, a one-way within-subjects BANOVA comparing the baseline, 0/6, and 6/6 update conditions was conducted to investigate whether the presentation of a second memory array affected recall error. Inconsistent with Chapter 5 and Kessler and colleagues (Experiment 1, 2015), substantial evidence against update condition impacting recall error was observed, $BF_{Update/Null} = 0.13$. Further follow-up tests showed evidence against a difference between the baseline and 6/6 condition, $BF_{10} = 0.32$. Consistent with Chapter 5, repeating the array twice in the 0/6 condition did not substantially impact recall error compared to the baseline (BF_{10} = 0.27), and 6/6 condition (BF_{10} = 0.29). Next, a 3 (partial-updates: 1/6, 2/6, 4/6) x 2 (probe location: repeated, updated) within-subjects BANOVA on recall error showed evidence against an effect of partial-update conditions, BF_{10} = 0.29, and the interaction, BF_{10} = 0.16, similar to the findings of Chapter 5 and Kessler and colleagues' findings. Consistent with the findings of Chapter 5, repeated probe locations did not yield lower recall error than updated locations, BF = 0.14. In line with the initial findings of Chapter 5, overall, the evidence for any effects of update condition or probe location on recall error were weak compared to Kessler and colleagues' original findings.

Table 13

Means (and Standard Deviations) of Recall Error as a Function of Update Condition.

	Update condition					
Probe location	Baseline	0/6	1/6	2/6	4/6	6/6
Repeated	83.00 (4.78)	82.69(9.14)	79.49(14.19)	83.49(12.51)	84.90 (12.93)	-
Updated	-	-	81.85 (11.69)	82.11 (14.38)	84.62 (11.32)	81.49 (10.65)

Following the analyses of Chapter 5, the models simulating temporary and permanent removal were compared by applying a widely applicable information criterion (WAIC; Gelman et al., 2014; Oberauer et al., 2017; Watanabe, 2021). The result of the model comparison was mixed (see Table 14). Model 1 best fit the 2/6 partial update and 6/6 full update conditions, whereas Model 2 best fit the 1/6 partial update and the 4/6 partial update condition. This may suggest that the models fitted do not adequately explain temporary and permanent removal of shape features in visual WM. Furthermore, this may indicate that updating shapes is more difficult than colour features. Further comparison between the recall error of Chapter 5 and the current chapter may be needed to understand whether this is the case.

Table 14

WAIC Values for Each Updating Condition (Partial and Full), Per Model Simulating Temporary and Permanent Removal.

Model	*Simulating*	*Updating Condition*	*WAIC*
1	Temporary Removal	1/6 Partial Update	33674.80
		2/6 Partial Update	31061.11
		4/6 Partial Update	36826.78
		6/6 Full Update	-4944.26
2	Permanent Removal	1/6 Partial Update	28296.88
		2/6 Partial Update	32885
		4/6 Partial Update	30029.27
		6/6 Full Update	-4252.35

Note. Lower WAIC values indicate better model fit.

The posterior parameter estimates of Model 1 are presented in Figure 13, and the differences of the partial-update conditions relative to the baseline, repeated (0/6) and full-update (6/6) conditions are presented in Table 15. Despite the findings of Chapter 5, it was observed that binding errors were similar in the model simulating temporary removal, with only a credible difference in the 1/6 condition (with an updated probe location): difference = 0.28 [0.04, 0.52]. While not credible (except for the 1/6 partial-update condition), the similarity suggests that binding errors occur regardless of the number of updates that are required and the probe location. Alongside the credible difference in the 1/6 partial-update condition, a credible difference in guessing in the 1/6 update condition (update probe location) was observed: difference = -0.35 [-0.61, -0.06]. No credible differences were observed in target memory or precision while comparing the partial-update conditions to the baseline condition, repeated, and full-update conditions.

Table 16 and Figure 14 show the results of Model 2. Similarly to Model 1, comparing the baseline to the partial-update conditions, credible effects of updating were observed in binding errors while 1/6 items were to-be-updated, and the probe location was update in nature (rather than repeat). Further credible effects were observed while comparing the 6/6 update condition to the partial-update conditions, but once again only while 1/6 items were updated. Credible effects were also observed in guessing while 1/6 items were updated. Given the similarity of the findings of Model 1 and Model 2, locally updating shape features may be more difficult than locally updating colour features. Further analyses would be required to determine whether this is the case.

Table 15

Model 1 Summary of Mean Differences [and 95% HDIs] in Each Memory Parameter of Partial-Update Conditions versus the Baseline, Repeated and Full-Update Conditions.

Comparison	Partial-Update Condition	Probe Location	*P*(Target)	*P*(Binding Error)	*P*(Guessing)	Precision
Baseline vs.	1/6	Repeat	0.06 [-0.03, 0.15]	-0.01 [-0.19, 0.16]	-0.05 [-0.27, 0.17]	-0.94 [-4.85, 3.15]
		Update	0.07 [-0.04, 0.17]	**0.28 [0.04, 0.52]**	**-0.35 [-0.61, -0.06]**	1.16 [-4.30, 8.75]
	2/6	Repeat	0.01 [-0.08, 0.10]	0.13 [-0.11, 0.43]	-0.14 [-0.49, 0.16]	-0.76 [-4.93, 3.78]
		Update	0.03 [-0.05, 0.11]	0.08 [-0.11, 0.28]	-0.12 [-0.34, 0.13]	2.14 [-4.06, 9.52]
	4/6	Repeat	-0.01 [-0.09, 0.07]	0.14 [-0.08, 0.38]	-0.14 [-0.41, 0.11]	2.06 [-3.92, 10.11]
		Update	-0.03 [-0.10, 0.05]	0.14 [-0.07, 0.37]	-0.11 [-0.38, 0.14]	2.59 [-3.28, 9.07]
Repeated (0/6) vs.	1/6	Repeat	0.04 [-0.06, 0.15]	-0.05 [-0.25, 0.14]	0.01 [-0.24, 0.27]	-3.89 [-10.32, 1.92]
		Update	0.05 [-0.07, 0.17]	0.24 [-0.04, 0.50]	-0.29 [-0.61, 0.02]	-1.79 [-9.67, 7.27]
	2/6	Repeat	-0.01 [-0.11, 0.09]	0.09 [-0.16, 0.38]	-0.08 [-0.42, 0.21]	-3.72 [-9.68, 2.04]
		Update	0.01 [-0.08, 0.11]	0.04 [-0.17, 0.30]	-0.06 [-0.35, 0.22]	-0.81 [-8.96, 7.79]
	4/6	Repeat	-0.03 [-0.12, 0.06]	0.10 [-0.16, 0.35]	-0.08 [-0.37, 0.22]	-0.90 [-8.77, 9.14]
		Update	-0.04 [-0.13, 0.05]	0.10 [-0.15, 0.32]	-0.06 [-0.33, 0.23]	-0.37 [-7.90, 7.35]
Full-Update (6/6) vs.	1/6	Repeat	0.04 [-0.07, 0.15]	-0.04 [-0.25, -0.17]	0.00 [-0.27, 0.24]	-4.37 [-11.72, 2.44]
		Update	0.05 [-0.07, 0.16]	-0.26 [-0.03, 0.52]	-0.30 [-0.61, 0.02]	-2.27 [-11.34, 7.20]
	2/6	Repeat	-0.01 [-0.11, 0.10]	-0.11 [-0.20, 0.43]	-0.10 [-0.48, 0.26]	-4.49 [-12.17, 2.93]
		Update	0.01 [-0.09, 0.10]	0.06 [-0.17, 0.29]	-0.07 [-0.34, 0.20]	-1.29 [-11.85, 7.75]
	4/6	Repeat	-0.03 [-0.12, 0.06]	0.12 [-0.13, 0.36]	-0.09 [-0.36, 0.19]	-1.37 [-10.29, 8.39]
		Update	-0.05 [-0.14, 0.04]	0.12 [-0.13, 0.35]	-0.07 [-0.36, 0.21]	-0.84 [-9.93, 7.52]

Note. Effects in boldface font indicate credible effects.

Figure 13

Posterior Parameter Estimates of Model 1 for the Probability of Target Memory (A), Memory Precision (B), Binding Errors (C) and Probability of Guessing (D). Larger Dark Circles Indicate Group Means, Smaller Faded Circles Indicate Individual Means, and the Error Bars Show the 95% HDIs of the Posterior.

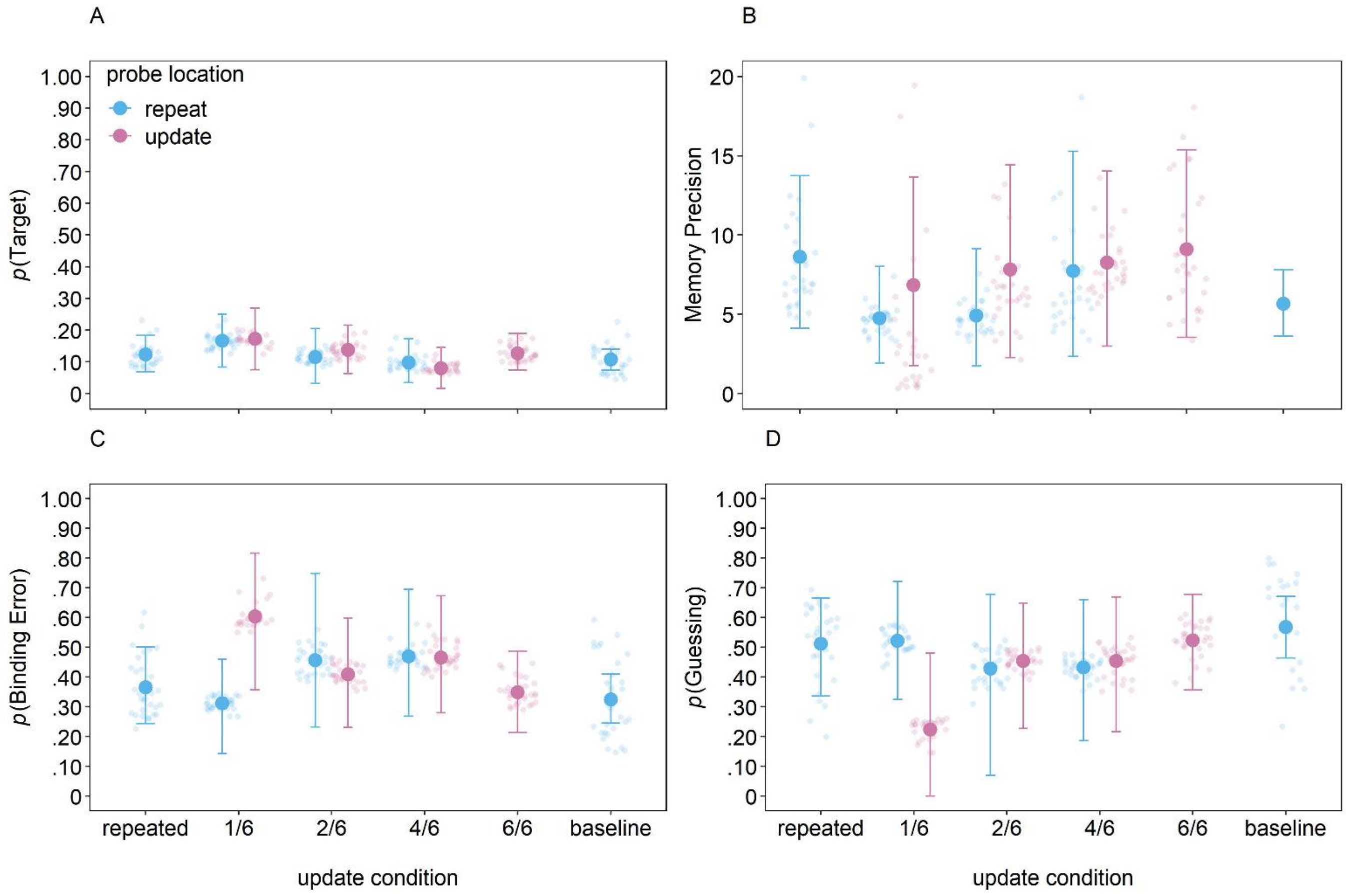

Table 16

Model 2 Summary of Mean Differences [and 95% HDIs] in Each Memory Parameter of Partial-Update Conditions versus the Full-Update (6/6) Condition.

Comparison	Partial-Update Condition	Probe Location	*P*(Target)	*P*(Binding Error)	*P*(Guessing)	Precision
Baseline vs.	1/6	Repeat	0.06 [-0.03, 0.16]	0.00 [-0.17, 0.20]	-0.06 [-0.31, 0.16]	-0.93 [-5.12, 3.29]
		Update	0.07 [-0.03, 0.17]	**0.27 [0.04, 0.50]**	**-0.34 [-0.59, -0.08]**	1.52 [-4.86, 9.00]
	2/6	Repeat	0.00 [-0.07, 0.07]	0.02 [-0.15, 0.20]	-0.02 [-0.23, 0.18]	0.21 [-4.23, 5.78]
		Update	0.05 [-0.04, 0.14]	0.08 [-0.11, 0.27]	-0.12 [-0.35, 0.11]	1.28 [-4.26, 8.03]
	4/6	Repeat	-0.01 [-0.09, 0.07]	0.08 [-0.14, 0.30]	-0.06 [-0.34, 0.18]	0.34 [-4.68, 6.29]
		Update	-0.02 [-0.09, 0.05]	0.09 [-0.07, 0.25]	-0.07 [-0.27, 0.11]	0.34 [-3.62, 5.07]
Repeated (0/6) vs.	1/6	Repeat	0.04 [-0.06, 0.16]	-0.04 [-0.25, 0.17]	0.00 [-0.28, 0.27]	-4.08 [10.36, 2.56]
		Update	0.06 [-0.05, 0.18]	0.23 [-0.02, 0.48]	**-0.29 [-0.58, -0.01]**	-1.63 [-9.13, 7.44]
	2/6	Repeat	-0.02 [-0.11, 0.07]	-0.02 [-0.23, 0.17]	0.04 [-0.21, 0.29]	-2.95 [-10.27, 3.59]
		Update	0.03 [-0.06, 0.14]	0.04 [-0.17, 0.25]	-0.07 [-0.32, 0.19]	-1.88 [-9.33, 6.21]
	4/6	Repeat	-0.03 [-0.12, 0.07]	0.04 [-0.19, 0.29]	-0.01 [-0.31, 0.27]	-2.82 [-9.86, 5.66]
		Update	-0.03 [-0.12, 0.05]	0.05 [-0.14, 0.25]	-0.01 [-0.25, 0.22]	-2.81 [-9.37, 3.11]
Full Update (6/6) vs.	1/6	Repeat	0.04 [-0.07, 0.16]	0.05 [-0.15, 0.24]	-0.09 [-0.35, 0.17]	-4.48 [-11.98, 2.85]
		Update	0.05 [-0.07, 0.16]	**0.32 [0.08, 0.56]**	**-0.37 [-0.64, -0.09]**	-2.03 [-12.06, 7.36]
	2/6	Repeat	-0.02 [-0.12, 0.07]	0.07 [-0.12, 0.28]	-0.05 [-0.30, 0.21]	-3.34 [-11.47, 3.76]
		Update	0.03 [-0.08, 0.14]	0.13 [-0.07, 0.34]	-0.16 [-0.41, 0.10]	-2.27 [-10.55, 6.87]
	4/6	Repeat	-0.03 [-0.13, 0.08]	0.13 [-0.10, 0.37]	-0.10 [-0.41, 0.19]	-3.21 [-11.55, 5.86]
		Update	-0.04 [-0.14, 0.04]	0.14 [-0.04, 0.32]	-0.10 [-0.33, 0.12]	-3.21 [-10.87, 3.40]

Note. Effects in boldface font indicate credible effects.

Figure 14

Posterior Parameter Estimates of Model 2 for the Probability of Target Memory (A), Memory Precision (B), Binding Errors (C) and Probability of Guessing (D). Larger Dark Circles Indicate Group Means, Smaller Faded Circles Indicate Individual Means, and the Error Bars Show the 95% HDIs of the Posterior.

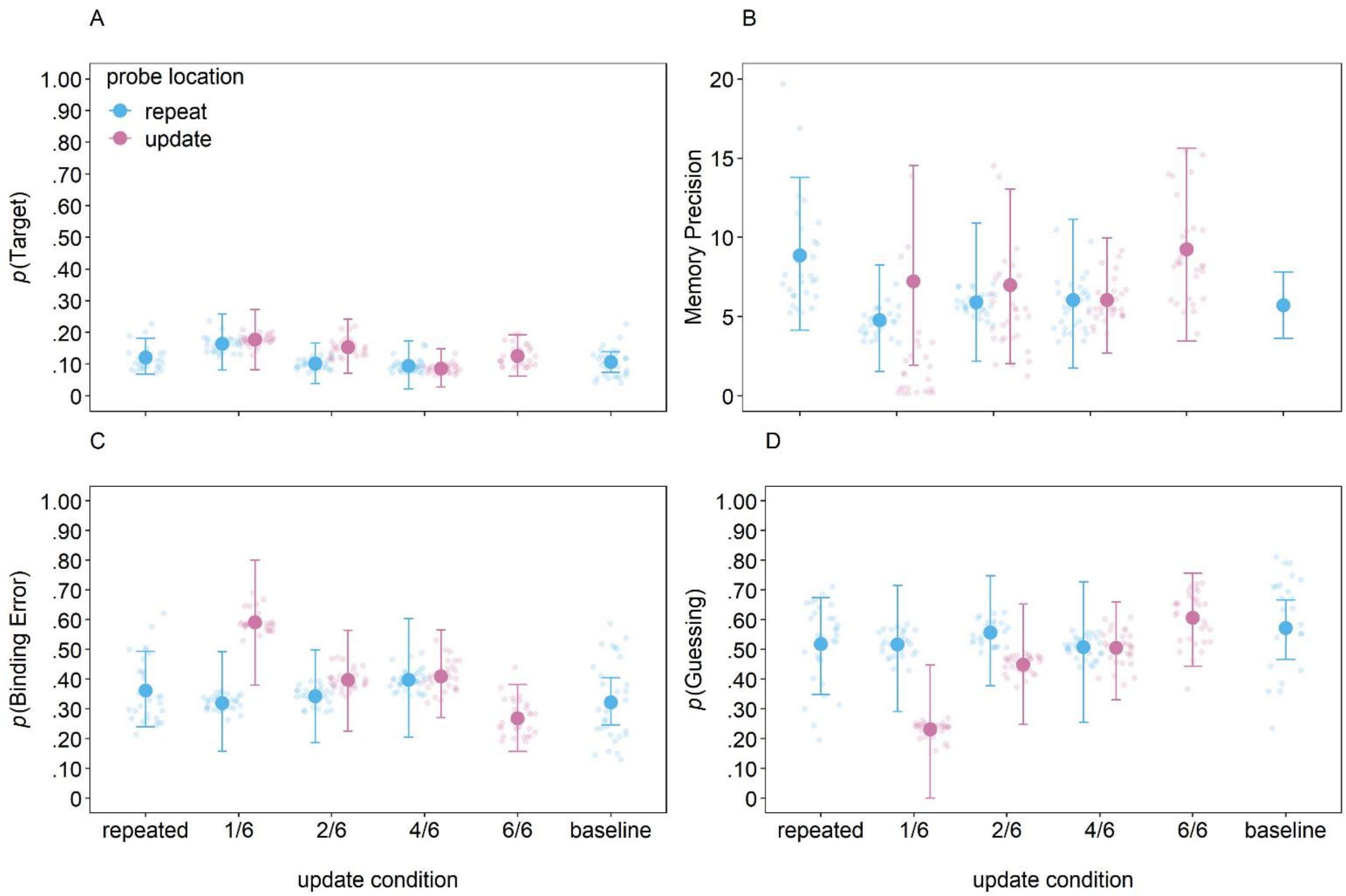

Exploratory Analyses

Given that recall error seemed higher overall in the current chapter compared to Chapter 5 (see Figure 15 for the plotted distribution of recall error), further analyses were conducted to understand whether updating shapes may be more difficult than colours overall. To this end, the previous analyses on recall error were repeated including experiment feature (colour or shape) as a between-subjects factor. First, the 3 (baseline, 0/6, and 6/6 update) x 2 (colour, shape) BANOVA showed that the full model including the interaction was the best model, BF_{10} = 40.24 x10^9. The model including the interaction was ambiguously preferred over the next-best model including only a main effect of feature type (BF = 1.04). Further paired Bayesian t-tests comparing recall error between colour and shape for each update condition indicated substantial and strong evidence that recall performance was worse for shapes than colours overall (see Table 17).

Figure 15

Distribution of Recall Error per Update Condition, Probe Location and Feature Type/Chapter (Colour/Chapter 5, Shape/Current Chapter).

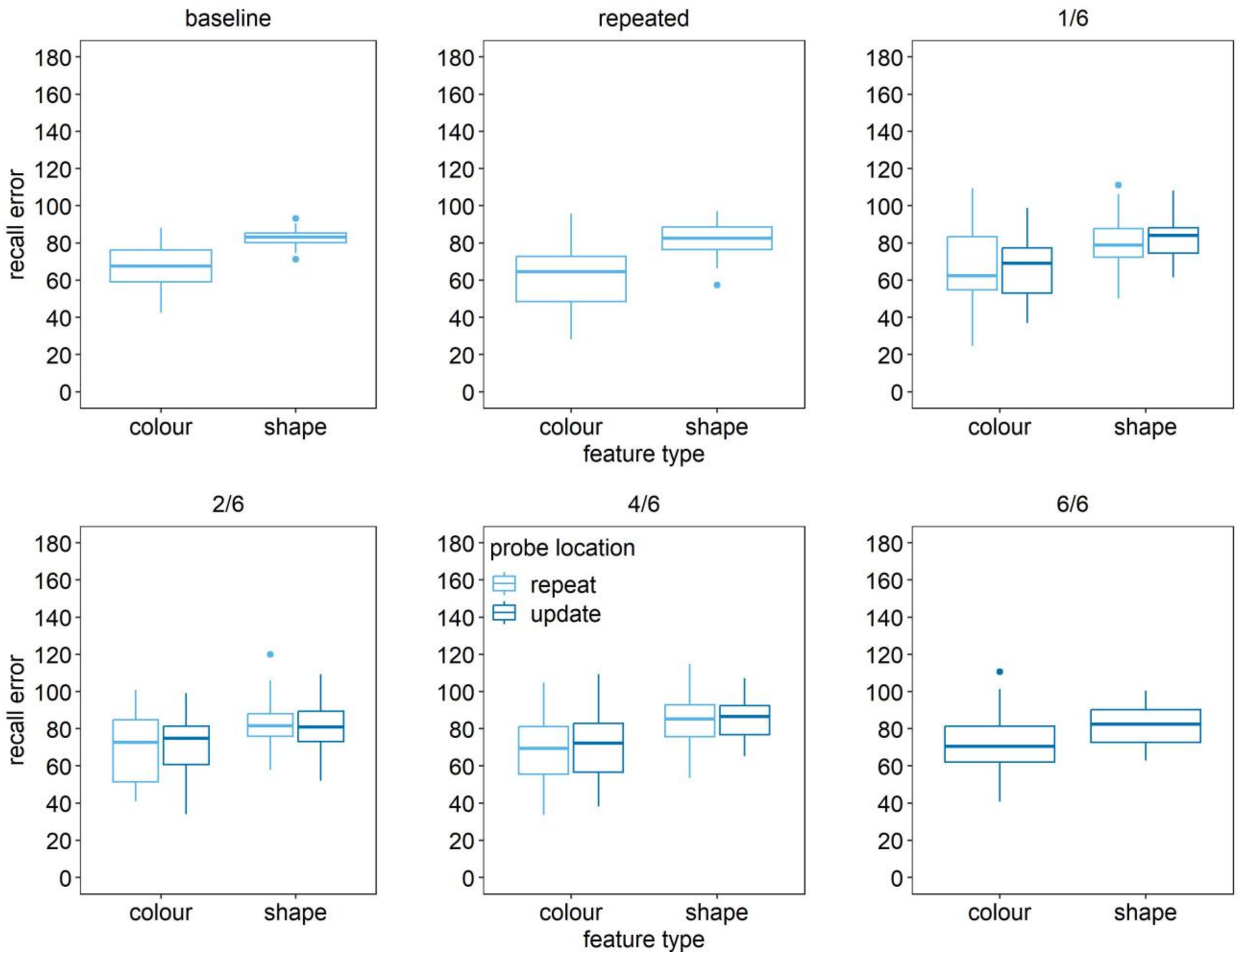

Table 17

Bayes Factors from the Bayesian T-Tests Comparing Recall Error Between Colour and Shape (Current Chapter) for each of Update Condition.

	Probe Location	
Update Condition	Repeat	Update
Baseline	74216.48	-
Repeated	4866.96	-
1/6	3.17	172.64
2/6	27.90	6.79
4/6	39.52	129.91
6/6	-	3.57

Furthermore, the 3 (1/6, 2/6, 4/6) x 2 (repeated, updated) x 2 (colour, shape) BANOVA showed that the model including only a main effect of the feature type (colour versus shape) was the best model, $BF_{10} = 4.41 \times 10^{16}$, which was substantially preferred over the next-best model including the main effect of the partial-update conditions and feature type (BF = 4.98). Further tests comparing the recall error between colour and shape for each of the partial-update conditions indicated substantial and very strong evidence that feature type impacts recall performance in the partial-update conditions, and irrespective of the probe locations (see Table 17). Taking these exploratory findings together, it can be argued that varying the feature type from colour to shape reduced task performance overall.

Discussion

In the current chapter, I investigated whether altering the to-be-updated feature-type from colour to shape would provide further evidence for the temporary removal of outdated information during local WM updating. The findings from the model comparison indicated that the model simulating temporary removal best fit the posterior distribution, relative to the model simulating permanent removal. Incongruent with the findings of Chapter 5, in Model 1, binding errors were similar regardless of the number of updates. This pattern of results was also observed while fitting Model 2 to the data. In both models fitted, target memory was low (see Tables 14 and 15), suggesting that updating shape features through the current paradigm may have proven to be difficult. Further between-subjects exploratory analyses revealed that recall error was indeed lower overall in the current experiment using shapes compared to Chapter 5 that used colours. This suggests that shapes are more difficult to maintain and update compared to colours, which is consistent with prior work (e.g., Allen et al., 2006). Nevertheless, that the model simulating temporary removal best fit the data compared to a model simulating permanent removal. It can therefore be argued that outdated shape information may still be held active in WM and outdated shape information is temporarily removed despite participants struggling with the current task.

One explanation for the overall poor performance in the current study's paradigm could be that not enough time was allowed for individuals to successfully store the shape features. Prior work has demonstrated that multiple colour features are stored faster than multiple shape features of real-world items (Brady et al., 2016). Brady and colleagues found that individuals can store multiple colour features at a low presentation time (less than 1000 ms) but storing multiple shape features of real-world objects requires longer presentation times (greater than 1000 ms). It could be possible that the fast presentation time of the shape features in the current study may have led to overall poor performance, as individuals were unable to store multiple presented shape features in the time given. To understand whether this is the case, further research investigating local WM updating of shape features (with presentation times greater than 100 ms) may be required. For example longer presentation times of 1000 ms (Chapters 2, 3, and 4) may be sufficient to allow for multiple shape features to be stored.

Moreover, Brady and colleagues suggested that WM storage capacity of multiple non-real-world items' shape feature, may be limited by the perceptual complexity of the shape, and its lack of conceptual meaning to an individual, which may have been the case here. Thus, further research into the removal of outdated shape feature information may be required. For example, it may be of interest to understand whether locally updating the shape feature of similar real-world objects relies on the temporary or permanent removal of outdated information. In day-to-day life, we encounter shape features of items that may be perceptually similar, but the feature belongs to a distinct item. For example, trainers vary in style (arguably, the trainer's shape feature), but the style belongs to a distinct pair of trainers. While browsing for a pair of trainers, one must rely on locally updating outdated shape features of in-memory trainers with new shape features. Additionally, one must also be able to compare outdated information with new information to complete one's goal (such as identifying what trainer style one prefers). Arguably, manipulating perceptually similar real-world items as to-be-updated may offer greater participant recall performance due to the participants' familiarity with real-world items. Therefore, it is of interest to understand whether one relies on the temporary removal of outdated real-world shape features that are perceptually similar, or whether the outdated information is permanently removed from WM.

In summary, I investigated whether the pattern of results observed in Chapter 5 (pertaining to colour) is also observable when shape features are to-be-updated. In line with the findings of Chapter 5, the comparison of the models simulating the temporary and permanent removal of outdated information suggested that the model simulating temporary removal best fit the data. However, overall performance was poor in the current chapter, and this may be due to participants not being able to easily store shape features when the presentation time is short. Further research is required to understand the removal of outdated shape features during local WM updating.

Chapter 7

General Discussion

The goal of this thesis was to investigate and understand how single features and feature bindings are prioritised during goal-oriented processing in visual WM. Firstly, I examined the role of attention during feature maintenance (Chapters 2 and 3). Secondly, I examined whether the prioritisation of location features is more beneficial than the prioritisation of colour features (Chapter 4). Thirdly, I examined whether outdated (irrelevant) features are removed from WM while new (updated) features are prioritised (Chapters 5 and 6). In this chapter, I will begin by summarising the findings of the current work. Secondly, I will address whether attentional resources are required to prioritise feature information during goal-orientated processing. Thirdly, I will discuss whether there are differences between the processing of unalike feature dimensions. Fourthly, I will discuss the processing of new features during WM updating. Finally, I will discuss what the implications of my findings are for how we understand the processing of single features and feature bindings in visual WM.

Summary of findings

In Chapter 2, by disrupting attention (Experiment 1) and guiding attention (Experiment 2), I investigated whether the role of attention during feature maintenance depends on the nature of the to-be-maintained feature (i.e., intrinsic, or extrinsic; Kessler et al., 2013; also see Chapter 2). The findings in Experiment 1 suggested that disrupting attentional resources did not increase intrinsic and extrinsic binding errors. However, disrupting central attention reduced target memory when the feature was extrinsic. Moreover, guiding attention in Experiment 2 via intrinsic and extrinsic retro-cues did not reduce binding errors relative to a no-cue condition. However, guiding attention improved target memory, irrespective of the nature of the retro-cue and tested feature type.

In Chapter 3, I aimed to replicate Experiment 1 of Chapter 2, with the exception that the set size of presented items was reduced from six items to four. The hypotheses remained the same for Chapter 3 as they were for Experiment 1 of Chapter 2. Despite reducing the number of to-be-maintained items to four (i.e., 4 items x 3 features (colour, shape, location) = 12 features), disrupting attention did not increase binding errors. Conversely to Experiment 1 of Chapter 2, no credible effects were observed in target memory, regardless of the disruption-type or the nature of the to-be-maintained feature.

Chapter 4 aimed to investigate whether there are differences between the prioritisation of colour and location features by adapting Experiment 2 of Chapter 2's paradigm. Similar to Experiment 2 of Chapter 2, no credible differences in binding errors were observed, regardless of the nature of the prioritised feature. Moreover, no credible differences were observed in binding errors between the no-cue conditions either, thus suggesting that there are no differences in prioritising location and colour features. However, target memory improved with the presence of a retro-cue, irrespective of its nature (colour-cue, or location-cue), thus confirming the findings of Experiment 2 in Chapter 2 that guiding internal attention improves target memory, irrespective of a feature's nature.

In Chapter 5, I examined whether prioritising new (updated) information relies on the temporary or permanent removal of outdated (irrelevant) information. I adapted a visual WM updating task to investigate whether applying two models simulating temporary and permanent removal would explain the processes involved with WM updating. The findings suggested that the model simulating temporary removal (i.e., including outdated and relevant information during model fit) best fit the data. Furthermore, binding errors increased as the number of to-be-updated items increased relative to the no-update baseline condition, particularly when more than one item was to-be-updated. This suggests that while outdated information is still available in WM, more opportunities arise for binding errors to occur. Conversely, the model that simulated permanent removal (i.e., only relevant information was included during model fit) provided a worse fit of the data, and binding errors were similar between the partial- and full-update conditions. This suggests that binding errors specifically arose from the outdated information.

In Chapter 6, I aimed to replicate and extend the findings of Chapter 5 by changing the feature dimension from coloured dots to unicoloured shapes. Unlike Chapter 5, binding errors were similar regardless of the partial- and full-update conditions and irrespective of the model fitted. However, the model simulating temporary removal best fit the data, in line with the findings of Chapter 5 suggesting that temporary removal may best explain how updating occurs in visual WM. The findings of Chapter 6, however, must be taken with a pinch of salt – I observed that overall performance was low compared to that of Chapter 5, thus suggesting that prioritising new (updated shape) features may be difficult in this experiment's paradigm.

Are attentional resources required for the processing of feature information?

Overall, the findings from Chapters 2-4 indicate that binding errors are not sensitive to disrupting or guiding attention. One explanation for this may be that binding errors are difficult to detect; this may be particularly true in the paradigms used in Chapters 2-4, wherein during recall, the probe-target types were randomly intermixed. The randomisation of the probe-target types during recall may have been challenging for the participants, and this may have led to guessing behaviours. However, guiding attention with retro-cues consistently improved target memory, regardless of the nature of the retro-cue or the target feature. This suggests that attentional resources may not aid WM maintenance by reducing the likelihood that binding errors occur during recall, but rather the findings suggest that attentional resources aid the likelihood that the target is kept active during goal-orientated processing. However, this raises the question of whether features are processed and prioritised in a domain-specific visual WM store, or whether the feature information is processed and prioritised in an attention-based domain-general framework.

As discussed in Chapter 1, the two principal perspectives suggest that maintenance in WM is either largely domain-specific (e.g., the multi-component model; Baddeley et al., 2020; Baddeley, 2000, 2012; Baddeley & Hitch, 1974) or domain-general (e.g., the embedded-processes model; Cowan, 1993, 1998, 1999; Cowan et al., 2020). While domain-specific accounts suggest that information is processed in their respective stores (visual information is processed in a visual store), the domain-general account suggests that attentional resources are similarly involved in maintaining goal-relevant information, regardless of its modality. Arguably, the findings from Chapters 2-4 suggest that attention may be required to keep features active in WM, irrespective of its nature (i.e., colour, shape, or location). I will now discuss these findings in more depth.

In Experiment 1 of Chapter 2, I only observed that disrupting attentional resources reduced target memory when the target feature was location-based and central attention was disrupted. Disrupting peripheral attention did not reduce target memory for colour-, shape-, or location-based features. This suggests that the maintenance of location-based features may be reliant on a general attentional resource and disrupting this resource, reduces one’s ability to retain the target location at a high quality. Conversely, the null peripheral disruption effects in target memory for colour and shape target features suggests that maintenance of colour and shape features may not rely on attention, but rather WM maintenance for these features does not require attention. As discussed in Chapters 1 and 2, Allen and colleagues (2006, 2012) posited that the binding of visuospatial information does not require attention. In line with this, the findings of Experiment 1 in Chapter 2 showed that individuals were able to retrieve target features while probed with its corresponding feature-type (i.e., colour probing location, colour probing shape, location probing colour, etc.) with a small likelihood that binding errors occurred. This may suggest that attentional resources are not required to bind visual information together. This was consistent even when reducing the set size from 6 to 4 in Chapter 3: Despite the reduction in set size and overall greater performance compared to Chapter 2, no credible differences in binding errors were observed between the disruption conditions and the baseline condition. Against the findings of Experiment 1 in Chapter 2, no credible differences were observed in target memory, irrespective of the nature of the probe-target combination and the disruption type. Given the opposite findings of Experiment 1 of Chapter 2 and Chapter 3, this suggests that attention may be necessary to maintain multiple feature bindings in visual WM when WM capacity has exceeded. As proposed in Chapter 3, one way to investigate whether this is the case is by calibrating the number of items (and features) that are presented to the participant. This would allow one to investigate whether disruption effects occur when a participant's WM capacity has exceeded.

But the findings of Experiment 1 of Chapter 2 and Chapter 3 raises the question: Do we observe a similar pattern of results while we guide internal attention, rather than disrupting attention? As explained earlier, Experiment 2 in Chapter 2 and Chapter 4 showed that retro-cues impacted target memory, and not binding errors, irrespective of the nature of the probe-target type. The findings of Experiment 2 in Chapter 2 contradict the findings of Experiment 1 – where I observed that attention does benefit the maintenance of to-be-retained information. This suggests that guiding attentional resources may have a stronger impact on keeping features active than the negative impacts of disrupting attention. With that said, it could be argued that the effects observed in Experiment 2 of Chapter 2 and Chapter 4 could be attributed to visual aftereffects rather than attention being guided internally towards the to-be-remembered feature. Arguably however this is not the case in the current work. Recent work guiding attention via retro-cues has shown a similar pattern of results as observed in Chapter 2 and 4, without the requirement of visual masking during the retention interval (e.g., Arnicane & Souza, 2021; Loaiza & Souza, 2019; Loaiza et al., 2023; Souza, 2016). Therefore, these findings concur with prior and current literature that suggests that attention can be focused on a subset of to-be-retained items, allowing for the memory trace of the to-be-retained item to be strengthened (Camos et al., 2018; Griffin & Nobre, 2003; Heuer & Schubö, 2016b; Ye et al., 2016).

The findings of the current work suggest that attentional resources may be an integral part of WM (e.g., Oberauer, 2019). However, as discussed in Chapter 1, it is still unclear how goal-relevant information is encoded and stored in domain-general versus domain-specific frameworks. It could however be argued that attentional resources (such as the focus of attention) "zoom-in" on the to-be-retained information, irrespective of its domain (i.e., visual, or verbal domain). This in turn activates the subset of LTM relevant to the domain that is to-be-prioritised. Incongruent with this conceptualisation of WM, Logie (1995) argued that visual information is stored and maintained in separate stores. It was argued that visual information is processed in a passive visual-cache, and spatial information is processed in a dynamic spatial-loop. Therefore, if visual information were to-be-stored in separate stores (such as the visual cache and spatial-loop), then I would have expected that guiding attention improved the target memory of the passively stored colour and shape features only. Against Logie's (1995) hypothesis, the findings of Experiment 2 in Chapter 2 reflect an account of WM that is not modal-specific (i.e., the dissociation of visual and spatial information), wherein guiding attention towards a to-be-prioritised feature strengthens the memory trace for the maintained in-memory items. It was observed that guiding attention improved target memory, irrespective of the nature of the retro-cue and the nature of the target feature. Therefore this suggests that single features and feature bindings may be stored in a domain-general manner, with attention strengthening the maintenance of visual information. However, this conclusion cannot be interpreted too strongly, as all the work in this thesis investigated the processing of a single domain (i.e., visual) rather than both verbal and visual modalities. Further work is required to determine whether verbal and visual information processing (and its cross-domain binding) can be explained by a domain-specific or domain-general perspective of WM. Therefore, considering the findings of the current work, it can be argued that processing information in the visual domain may be benefitted by attention.

Processing dissimilar feature dimensions

Prior research has suggested that location features may be privileged in WM and in perception, additionally attention aids the binding of non-location features to location features (Johnson & Pashler, 1990; Rajsic & Wilson, 2014; Schneegans & Bays, 2017; Treisman & Gelade, 1980). The findings from Experiment 2 in Chapter 2 suggested that there are no differences between the maintenance of feature dimensions, and attention aids the prioritisation of to-be-maintained features irrespective of their type. Chapter 4's paradigm was adapted from Experiment 2 in Chapter 2 to determine whether there are differences between maintaining non-location and location features, and whether there are differences between the prioritisation of non-location and location features. As summarised earlier, no credible differences were observed between the baseline conditions of the colour and location feature types (i.e., uncued retention interval). This may suggest that the processing of colour and location features is automatic (Allen et al., 2006, 2012). This finding disagrees with the feature integration hypothesis of visual perception (FIT; Treisman & Gelade, 1980; Wheeler & Treisman, 2002). As discussed in Chapters 1 and 4, FIT proposes that single features are stored in their respective feature dimensions and attentional resources are required to bind non-location features to their respective location feature. While FIT is a hypothesis concerned with the initial processing of visual features in perception, arguably its importance must be considered during visual WM due to its overlap with the role of attention during goal-orientated processing. It can be argued that similarities can be drawn between FIT hypothesis of visual perception and Logie's (1995) explanation of visual processing in WM (i.e., the inner scribe and visual cache). For example, in both hypotheses, location features appear to be critical to the processing of non-location features; however, in FIT, attentional resources aid the binding and maintenance of non-location features and their respective locations during immediate perception. Therefore, it is critical to understand whether in WM, there are differences between the processing of non-location and location features during attentional prioritisation.

To answer this question, Chapter 4 also investigated whether there are differences between guiding internal attention towards a to-be-prioritised non-location and location feature. It was hypothesised that guiding internal attention via a location retro-cue would reduce colour binding errors, relative to a no-cue baseline condition, due to the privileged state of location features during the attentional maintenance of visual features in their respective feature dimensions. Furthermore, it was hypothesised that guiding internal attention via a colour retro-cue should not reduce location binding errors, relative to a no-cue baseline condition. Finally it was hypothesised that if non-location features are not bound to their respective locations and attentional resources do not selectively maintain location features, then binding errors should be similar, irrespective of the cue condition (i.e., none, colour-cue, and location-cue). Chapter 4's findings suggest that there are no credible differences in binding errors irrespective of the target feature, thus confirming the third hypothesis. Although there were no credible differences in binding errors, similar to Experiment 2 in Chapter 2, credible improvements were observed in target memory and guessing irrespective of the nature of the retro-cue and the target feature type.

Against FIT, the findings of Chapter 4 suggest that attention is not selectively oriented towards locations during visual WM maintenance. Rather, attention may be thought of as a global process that strengthens in-memory representations regardless of their nature (e.g., colour or location). This concurs with my previous section: Attention is not selective in nature. Moreover, attention is able to "zoom-in" or focus on a feature, and thus, improving the quality of in-memory representations. Furthermore, the findings of Chapter 4 highlight that visual WM may rely on the maintenance of bound representations. This suggests that the binding of feature information is critical to allow for task-relevant information to-be-kept active. Therefore, it can be suggested that WM may encode and store features of an item as bound representations, thus concurring with the object-based view of visual WM (i.e., Luck & Vogel, 1997; see Luck & Vogel, 2013 for a review; Vogel et al., 2001). However WM may be reliant on attention to maintain the bindings and for high quality representations to-be-kept active.

Processing new features during WM updating

In Chapters 2-4, I examined whether attention is required to prioritise to-be-retained information, however, what happens to outdated (irrelevant) information when we are required to prioritise new (updated) information? As discussed in Chapters 1 and 5, WM updating is one's ability to alter and manipulate in-memory information (Lewis-Peacock et al., 2018). Although WM updating research has investigated how information may be updated, little is known about the processes that underly the prioritisation of the new information, and the removal of the outdated information. Therefore, in Chapters 5 and 6, I examined whether local WM updating relies on the temporary or permanent removal of outdated information. Temporary and permanent removal were simulated and compared by fitting the data to two hierarchical Bayesian three-parameter mixture models (Oberauer et al., 2017) that either included both outdated and new information (Model 1, temporary removal) or only included new information during model fitting (Model 2, permanent removal). The findings in both Chapter 5 and 6 suggested that the model simulating temporary removal best fit the data, with binding errors increasing with increased items to update relative to a no-update baseline condition in Chapter 5 (but not Chapter 6). This suggests that prioritising new (updated) colours during WM updating may be susceptible to binding errors. Moreover, binding errors occurred regardless of whether the probe was repeated or updated. As alluded to in the prior section, it may be possible that visual WM stores and processes information as bound representations. As discussed in Chapter 5, it was argued that the increase in binding errors while updating multiple colours, and while outdated information is still accessible, suggests that outdated information may still be bound to the retained information. Therefore during WM updating and maintenance, it could be possible that both task-relevant bindings and outdated bindings may be kept active. This would suggest that additional resources may be required to "hone onto" goal-relevant bindings, and in this way, irrelevant bindings may be inhibited. But the inhibited irrelevant bindings interfere with the goal-relevant items, particularly when multiple irrelevant bindings are accessible. This is congruent with Palladino and colleagues' (2001) suggestion that during WM updating, outdated information may be suppressed. Arguably, this suggests that the proposed resource is limited and may only hold one binding active at a time (similarly to the model proposed by Oberauer, 2002).

Furthermore, the findings from Chapter 5 concur with the local WM updating hypothesis (Kessler et al., 2015; Kessler & Meiran, 2008). The fact that the pattern of binding errors did not occur for shapes in Chapter 6 as it did for colours in Chapter 5 suggests that the processing of colour and shape features differ during WM updating. However, it may be the case that colours can just be processed more quickly and accurately than shapes, and thus the paradigm was not optimal for this reason (see Chapter 6). Further research is required to understand whether the processing of unrelated features differ during WM updating.

Processing and prioritising item representations in visual WM

In the above sections, I have made the case that visual WM may maintain visual information in the form of bound representations, and that additional resources are crucial for keeping goal-relevant information active. Disrupting attention (Experiment 1 of Chapter 2) had no effect on binding errors, but disrupting central attention reduced target memory while the target feature was location-based. Similarly to Experiment 1 of Chapter 2, disrupting attention (Chapter 3) had no effect on binding errors, but unlike Experiment 1 of Chapter 2, target memory was unaffected by the disruption-types. However, further analyses revealed that reducing the set size in Chapter 3 (from 6 items to 4) improved task performance. Guiding attention via retro-cues (Experiment 2 of Chapter 2, and Chapter 4) had no effect on binding errors, but retro-cues improved target memory, irrespective of its nature or the retro-cue or the target feature. Taken together, these results can be interpreted to suggest that the role of attention in visual WM is to keep goal-relevant bindings active. Moreover, the findings from Chapter 3 suggest that when the

WM capacity limit has not been exceeded (such as the fixed-item capacity proposed by Cowan, 2001, and Zhang and Luck, 2008), visual WM may be less reliant on drawing on attention to keep task-relevant bindings active. Following this logic, this suggests that attention may be a flexible resource that is only drawn upon when it is required. Therefore attention may be governed by a hypothetical component in WM, such as the central executive in the multi-component model (Baddeley and Hitch, 1974; see Baddeley 2012 for a review). Congruent with the domain-general conceptualisation of WM (such as, Barrouillet & Camos, 2010; Cowan, 1999; Oberauer, 2002; Vergauwe et al., 2009), the findings of Chapters 2 – 4 indicate that features of an item are stored in a holistic manner, such that their bindings are held active, rather than in a dichotic manner (i.e., visual information such as colour and shape, and spatial-based information, such as location are stored separately within a visual WM store; Baddeley, 2012; Logie, 1995). However, unlike the domain-general view, attention may be moderated by a hypothetical component, such as the multi-component's model central executive (Baddeley & Hitch, 1974).

Moreover, prior work has argued that item features are stored independently within their respective feature dimensions and attention aids the binding of features from within their feature dimensions (Treisman & Gelade, 1980; Wheeler & Treisman, 2002). Treisman and Gelade (1980) proposed that attention is selectively orientated towards location-based features. The findings from Chapters 2–4 suggest that this may not be the case.
As discussed earlier, while disrupting central attention in Experiment 1 of Chapter 2, it was observed that target memory for location features was reduced, but target memory for colour and shape features was unaffected. However, I was unable to replicate this finding while reducing the set size in Chapter 3. Conversely, it was observed in Experiment 2 of Chapter 2 and Chapter 4 that guiding attention improved target memory, irrespective of the nature of the retro-cue and target

feature-type. These findings are incongruent with FIT (Treisman & Gelade, 1980; Wheeler & Treisman, 2002). While the finding of Experiment 1 of Chapter 2 may agree with FIT, I did not observe the same pattern of results in Chapter 3. Arguably, if attention is selectively orientated towards location-based features, then in Chapter 3, disrupting central attention should have reduced target memory while the target feature was location-based. Similarly, in Experiment 2 of Chapter 2 and Experiment 4, I should have observed that target memory would only have improved while attention was guided via a location-based retro-cue, rather than target memory improving irrespective of the retro-cue's nature.

To summarise, it can be argued that the processing of visual information in WM may be conceptualised as domain-general and item features are not stored in feature-specific stores, rather visual WM may store the bindings of item representations (i.e., colour-shape, colour-location, or colour-shape-location bindings). Moreover, attention benefits the likelihood of target bindings being kept active in WM without bias towards a feature-type. However, attention is arguably a dynamic resource that can be called upon when needed (such as in Chapters 2 and 4, but not in Chapter 3). This however suggests that attention may be directed, or governed, by a component within WM. It can be posited that rather than attention being governed by a component within WM, the flexible nature of attention may allow for it to be "switched on" or "switched off" as necessary. Therefore, if attention is not needed to prioritise goal-relevant bindings in WM, then it does not need to be active. This view agrees with more recent work that has argued that during WM in the human brain, networks associated with attention are recruited when WM is required to prioritise goal-relevant information (see Brissenden & Somers, 2019 for a review).

Arguably, human cognition is a collection of multiple cognitive processes, such as memory, attention, decision making, problem solving, comprehension, reasoning, et cetera. It can therefore be hypothesised that when WM does not require attention to keep information active, and attention is switched "off", then attention can be directed towards the other aforementioned processes. If this is the case, then this suggests that WM can call upon other processes to aid information processing, and other cognitive processes can call upon WM. For example, while making a decision about a future goal, one must be able to hold multiple pieces of information active to make an informed decision. Therefore, it can be argued that a collection of cognitive processes is required to achieve one's goal, namely: WM (to hold multiple pieces of information active for a short period of time), attention (to focus on critical information within WM, and to switch between the information), and decision making (to weigh the information held active in WM in relation to one's goal). Taking this into account, it can be posited that cognitive processes are flexible, or dynamic.

While visual information may be represented as feature bindings in WM, and attention may be thought of as flexible, do we observe that irrelevant (outdated) bindings are no longer active in WM? The model comparison findings of Chapters 5 and 6 suggested that the model that best fit the data was the model simulating temporary removal. These findings indicate that during WM updating, irrelevant bindings may still be accessible. As discussed earlier, this suggests that not only are updated-retained bindings accessible in WM, outdated-retained bindings are also present. In Chapter 5, it was observed that binding errors increased with the increased number of items that were to-be-updated (2/6, 4/6, and 6/6), and irrespective of the probe location. Arguably, this indicates that keeping multiple outdated-retained bindings in WM may interfere with the prioritisation of updated-retained bindings. It may be posited that retaining updated-retained bindings is reliant on a flexible resource but once the outdated-retained bindings outweigh the updated-retained bindings, the flexible resource is "switched off" and more opportunities for errors to occur arises. However, this pattern of results was not observed in Chapter 6. In Chapter 6, it was observed that binding errors were similar across updating conditions, and irrespective of the probe location. I interpret this to suggest that updating shape features is more difficult than colour features. Further analyses comparing the recall error of Chapter 5 and Chapter 6 revealed this to be the case. Against my earlier discussion, that the processing of features is not reliant on independent storage of features in individual feature stores, this suggests that WM updating of item features differs between feature-types.

One way to investigate whether there are differences between the updating of feature-types in WM may be through a whole-report paradigm, in which the paradigm used in Chapters 5 and 6 is adapted: Coloured shapes are presented during the encoding phase and during the update phase, either the colour, shape or both are updated. During retrieval, participants would be required to recreate both the colour and the shape feature using two colour and shape wheels, rather than the single feature wheels used in Chapters 5 and 6. This paradigm would allow one to investigate whether there are differences between the updating of dissimilar (but bound) feature-types. Moreover, this paradigm may elude as to whether outdated-retained bindings interfere with updated-retained bindings. Additionally this may allow one to investigate whether the proposed flexible resource may be gated by a prioritisation capacity, in that, the flexible resource is limited by the number of outdated-retained bindings outweighing the number of updated-retained bindings and therefore the flexible resource is "switched off". Therefore, further research is required to understand how visual information is prioritised for goal-orientated processing during WM updating.

Considering all the work in this thesis, it can be suggested that visual WM processes information in the form of bound representations (feature bindings rather than single features), and additional resources such as a flexible attentional resource benefits the prioritisation and maintenance of task-relevant bindings. As such, the findings from the all the work in this thesis suggest that prioritising correct bindings is critical for accurate memory retrieval. Furthermore, it can be posited that a flexible resource (such as attention) may be "switched on/off". This hypothesised resource however may be subject to interference when the number of irrelevant bindings outweigh the number of relevant bindings.

Therefore it can be argued that the processing (and subsequent prioritisation) of item representations may rely on an extended object-based view, in that, items are represented in a bound representation, but additional resources improve the quality of prioritised feature bindings. This draws parallels with Zhang and Luck's (2008) slots + resource hypothesis. To reiterate, Zhang and Luck (2008) proposed that items are processed and stored as bound representations in fixed-item capacity "slots". The slots + resource hypothesis suggests that bound representations are stored in fixed-item capacity slots, but an additional resource can be called upon to improve the precision of a single binding. Additionally, Zhang and Luck proposed that all of the slots could be allocated towards a single binding, and the binding is distributed across the slots (slots + averaging hypothesis).

I interpret the findings from this thesis to suggest that the processing of visual information occurs through an extended version of the slots + resource hypothesis. As discussed previously, my findings suggest that attention may be flexible. In Experiment 1 of Chapter 2 it was observed that disrupting central attention reduced target memory but reducing the set size in Chapter 3 did not replicate this result. This suggests that visual WM capacity may be limited by the number of items or features that it can keep active, however attention may be required to keep multiple items or features active. Therefore, it is posited that attention can be "switched off" when it is not required to hold multiple items active when WM capacity has not been exceeded. Furthermore, the findings of Experiment 2 of Chapter 2 and Chapter 4 suggest that WM processes visual information as bindings, rather than single features. It was observed that guiding attention via retro-cues improved target memory, irrespective of the nature of the retro-cue or the target feature. Thus indicating that attention can be “switched on”, and focused onto a single representation, to improve the quality of a target feature binding. Moreover, the observed retro-cue benefit in target memory in Experiment 2 of Chapter 2 and Chapter 4, irrespective of the nature of the retro-cue and target feature-type suggests that attention is not biased towards a single feature. Thus, drawing upon attention during WM maintenance improves the quality of the prioritised feature binding. In addition, the findings from Chapter 5 indicated that the model that best fit the data was a model simulating temporary removal. Furthermore, it was observed in Chapter 5 that binding errors increased with the increased number of updates. This suggests that outdated information may interfere with relevant information. As such, outdated bindings may limit the prioritisation of new (updated) bindings, causing opportunities for errors to occur. Congruent with Chapter 5, Chapter 6’s model comparison indicated that the model that best fit the data was the model simulating temporary removal. Incongruent however, it was observed that binding errors were similar irrespective of the update condition, probe location and model fit. Further analyses indicated the incongruent findings may be due to updating shape features being more difficult than updating colour features.

Taking these findings together, it can be posited that feature bindings may be stored in a similar fashion as Zhang and Luck’s (2008) slots + resource hypothesis. However, unlike Zhang and Luck’s (2008) hypothesis, the findings from this thesis indicate that the additional resources may be attentional based and do not need to be activated at all times. Furthermore, the findings from this thesis suggest that the attentional resource can selectively prioritise a binding to strengthen the quality of the representation. Arguably this is similar to the slots + averaging hypothesis (Zhang & Luck, 2008). However, the slots + averaging hypothesis does not account for attentional prioritisation. Moreover, the additional resources may be subject to interference when the number of outdated bindings outweigh the relevant bindings, thus causing opportunities for errors to occur. Therefore, the current findings of this thesis suggest that the processing of visual information in WM is object-based, but additional resources (such as attention) can be called upon when required.

Limitations of current work

It can be argued that a limitation of this thesis may have been that online data collection was used throughout the current work. At the time of starting the thesis, a global pandemic prevented in-person data collection, and most researchers moved to online platforms. One may argue that it is difficult to observe psychological phenomena in an online setting. On the other hand, it can be argued that online data collection may provide for a "truer" representation of one's cognition as it is closer to everyday life than inviting participants to complete a high intensity study in a laboratory setting. Recent research comparing online, and in-person methodologies has shown that similar memory performance can be observed between online and in-person methodologies (Elliott et al., 2022). Thus, it can be argued that the performance of online data collection methods is just as adequate as the performance of in-person data collection methods. However, further research using my paradigm in a laboratory setting may be required to determine whether differences would be observed in the current work.

A second limitation of the current work may be that the aims of this thesis were only examined by means of continuous-report paradigms. However, as discussed in Chapters 1, 2 and 5, continuous-report paradigms offer an insight into the underpinnings of WM recall. Arguably, it may have been beneficial to the current thesis to also examine my hypotheses using alternative paradigms and alternative computational models. For example, the current work could have examined the prioritisation of feature information through a 3 alternative forced choice (3-AFC) paradigm, and the recall responses could have been fit to a multinomial processing tree (MPT) model (for a review, see Erdfelder et al., 2009). Using these methods would allow the determination of whether the findings observed in the current work may also be observed using alternative research methods. This would strengthen the case of the current work: Prioritisation of feature information during goal-orientated processing relies on the initial processing of an item's features.

It can be argued that a third limitation of the current work is that it is unclear as to whether individuals truly relied on their visual WM during the experiments in the current work. It could be argued that individuals may have relied on a mixture of visual and verbal encoding of the visually presented information. For example, in Experiments in Chapters 2-4, the longer presentation times of the to-be-retained information (1000 ms) may have enabled participants to verbally encode features of the multi-feature items (such as the items' colour). If this had been the case, however, I would have observed that precision was consistently higher in one feature dimension than in the others. This was not the case. Furthermore, the short presentation time of the to-be-retained items in Chapters 5 and 6 (100 ms) may not have been enough time to verbally encode all the presented feature information, particularly during the updating phase (where one or more than one item was to-be-updated). Therefore, while this thesis may have methodological limitations, it can be argued that all the findings of this thesis offer an insight into how items are represented in visual WM.

Future directions

As discussed in the previous section, it is unclear in the current works' experimental paradigms whether individuals exclusively encoded the visually presented items using visual WM, or whether a mixture of visual and verbal encoding was used. It therefore may be of interest to future researchers to adapt the current works' paradigms to include articulatory suppression throughout the encoding phases and retention intervals. This would allow researchers to identify whether processing of feature bindings occurs through a mixture of encoding the verbal and visual encoding, or whether feature bindings are processed by the visual domain only.

As discussed previously, it can be argued that the lack of similarity between the findings of Chapters 5 and 6 suggest that while outdated information is temporarily removed, colour and shape features are processed in a different manner during WM updating. Further research is required to explore this possibility. For example, it may be of interest for future researchers to investigate whether there are distinct differences between the processing of colour and shape features during WM updating by adapting the WM updating paradigm used in Chapters 5 and 6. Presentation times could be extended from 100 ms to 500 ms to allow for adequate time for the encoding of shape features. Future researchers may also present an array of coloured shapes during the encoding phase, rather than an array of coloured dots or unicoloured shapes. This would allow additional conditions to be set, in which either the colours of the items, or the shape of the items update randomly. Using this approach would allow exploration of whether there are distinct differences between WM updating of interdependent feature dimensions.

It may also be of interest for future researchers to apply neuroscientific methods towards the goal of understanding the prioritisation of single features and feature bindings during goal-orientated processing. For example, there is a trend emerging in cognitive neuroscience that investigates the connectivity between networks during cognition, using mathematical models such as graph theory (for a review, see He & Evans, 2010). It may be of interest to future researchers to apply these neuroscientific methods to the paradigms used in the current work (particularly in Chapters 2-4) to investigate whether there are differences in network connectivity between feature types, or whether the network connectivity pattern reflects that items are represented in a bound state. Arguably, the current work may offer a steppingstone for future research that is interested in the disentangling the processing of features in visual WM.

Conclusion

In conclusion, the aim of this thesis was to examine the processing of single features and feature bindings. The current work examined whether attention is required to maintain and prioritise goal-relevant features in visual WM. It was demonstrated that guiding attention benefits the prioritisation and maintenance of goal-relevant features. Furthermore, this thesis examined whether prioritising new features during WM relies on the temporary or permanent removal of outdated features. Overall, the current work suggests that outdated information is temporarily removed from WM and still accessible. The current work also highlights that updating shape features is more difficult than updating colour features in WM. Taking all the findings of this thesis into consideration, I argue that the processing of visual information can be thought of as relying on slot-like encoding and storage of feature bindings. Additionally, I argue that the maintenance of feature bindings is benefitted by a flexible attention-based resource that can be “switched on”, depending on whether it is required. Furthermore, the proposed resource may be subject to interference when the number of irrelevant (outdated) bindings outweighs the number of relevant bindings. Finally, I argue that the findings of this thesis suggest that items are initially processed in an object-based manner in WM (similarly to the slots + resource hypothesis), but additional resources improve the prioritisation of goal-relevant feature bindings. Therefore this thesis offers an insight into the processing of visual information in WM, but further research is required to understand whether the proposed explanation of visual information processing can be generalised.

Appendix A

Mixture modelling

As previously discussed, Zhang and Luck (2008) investigated WM capacity of visual items. They investigated whether visual items are stored in "slots" by fitting a two-parameter mixture model to the observed responses in a continuous-report paradigm, denoted as:

$$(1)\ p(\hat{}\theta)\ = (1-\gamma)\phi_{\sigma}(\hat{}\theta-\theta)+\gamma\frac{1}{2\pi}$$

Where θ is the target feature value in radians (e.g., colour), $\hat{}\theta$ is the reported feature value, and γ is the ratio of trials on which the participant responds at random. ϕ_{σ} represents the circular Gaussian distribution (von Mises distribution) with a mean of 0 and a standard deviation σ. Zhang and Luck (2008) proposed that there are two sources of error on each trial, namely the Gaussian variability in target memory (P_m), and a fixed probability of random guessing (P_{guess}). Thus, the precision (*k*) of a retrieved feature (such as colour) may be due to a mixture of target memory and random guessing. While Zhang and Luck's two parameter mixture model is fundamental in understanding the underlying memory parameters during feature retrieval, the model fails to address whether binding errors occur during feature retrieval.

The three-parameter mixture model proposed by Bays and colleagues (2009) adds a third component to understand the underlying memory parameters involved with feature retrieval, the likelihood of recalling a presented but non-target feature (i.e., a swap/binding error. For example it is possible that during visual WM maintenance, the colour of an item in position *x* is misbound with the colour of an item in position *y*. As such, the colour of the item in position *y* is retrieved rather than the target colour for position *x*. Bays and colleagues (2009) formalised their model as:

$$(2)\ p(\hat{}\theta) = (1-\gamma)\phi_\sigma(\hat{}\theta - \theta) + \gamma\frac{1}{2\pi} + \beta\frac{1}{m}\sum_i^m \phi_\sigma(\hat{}\theta - \theta_i^*)$$

Where β is the probability of misremembering the target location and $\{\theta_1^*, \theta_2^*, \ldots\ \theta_m^*\}$ are the feature values of the non-target items, m. Unlike Zhang and Luck's (2008) two-parameter mixture model (Equation 1), k of a retrieved feature may be due to a mixture of P_m, P_{guess} and P_{nt} (i.e., probability of recalling a non-target feature). While Bays and colleagues' (2009) three-parameter mixture model offers further insight into the underlying memory parameters during feature retrieval (through the parametrisation of misbinding), one caveat of the three-parameter (and two-parameter) mixture model is that it is limited by the sum of the probability parameters to into one parameter. Therefore caution must be taken to interpret whether the observed memory parameters reflect feature retrieval. One way in which one could achieve better interpretation of a mixture models' memory parameters is through a hierarchical modelling approach.

Hierarchical Bayesian three-parameter mixture model

Hierarchical modelling offers … (see Oberauer et al., 2017 for a discussion of hierarchical modelling). Oberauer and colleagues (2017) argued that applying a hierarchical modelling approach to the three-parameter mixture model (Bays et al., 2009) allows researchers to measure the memory parameters (k, P_m, P_{guess} and P_{nt}) on an individual level (i.e., one person's retrieval probability, rather than a group mean retrieval probability) and on a group level (i.e., the mean retrieval probability), and the group level's standard deviation (allowing for individual differences). Therefore, Oberauer and colleagues (2017) reparamertised the three-parameter mixture model (Bays et al., 2009) to account for individual differences during retrieval in a hierarchical Bayesian three-parameter mixture model:

$$(3)\ Pmem = Pm \cdot Pt$$

$$Pnt = Pm \cdot (1 - Pt)$$

P_m is the probability that a participants' response comes from any of the presented items, including the target item and the non-target items. P_t is the probability that the response mirrors the targets' feature (e.g., colour). The hierarchical Bayesian three-parameter mixture model predicts that each response during retrieval, $\hat{X_{i,j}}$,1 of participant *j* in trial *i* is distributed according to a von Mises distribution, wherein $m_{i,j}$ is the mean of the von Mises distribution and it is the true feature of the item from which the response derives:

$$(4)\ x_{i,j} \sim VM(m_{i,j}; z_{i,j} \cdot k_j)$$

The item determining the response is sampled from the array of presented items according to a categorical distribution with probability $P_{i,j}$:

$$(5)\ m_{i,j} = M_{i,j}\left(item_{i,j}\right)$$

$$item_{i,j} \sim Cat\left(P_{i,j}\right)$$

$$P_{i,j}(1) = Pt_j$$

$$P_{i,j}(2:n) = \frac{(1 - Pt_j)}{n - 1}$$

$M_{i,j}$ is a vector of memory features in an array *i* presented to participant *j*. The first element is the target feature and $P_{i,j}$ is a vector of probabilities of the responses of each of the features within an array of presented items. The response is related to the sampled memory feature $m_{i,j}$, with probability P_{mj}, whereas with probability 1 - Pm_j the retrieval response is drawn from a uniform distribution. The precision parameter *k* = 0 is a uniform distribution under the von Mises distribution. Therefore, Oberauer and colleagues (2017) multiplied k_j with a binary variable $z_{i,j}$ drawn from a Bernoulli distrubtion:

$$(6)\ z_{i,j} \sim Bern(Pm_j)$$

When $z_{i,j}$ = 1, the retrieval response is drawn from a von Mises centred on $m_{i,j}$ whereas when $z_{i,j} = 0$ the retrieval response is drawn from a uniform distribution. Therefore, the memory parameters for each participant (i.e., target memory, target feature, and precision; Pm_j, Pt_j, and k_j) are drawn from distributions with their own variability across participants. Thus allowing one to understand the individual differences of the memory parameters through:

$$(7)\ Pm_j \sim \beta(A_{Pm}, B_{Pm})$$

$$Pt_j \sim \beta(A_{Pt}, B_{Pt})$$

$$k_j \sim \gamma(S_k, R_k)$$

Furthermore, to analyse the memory parameters on a group level, Oberauer and colleagues (2017) formalised the priors that should be fit to the retrieval responses:

$$(8)\ A_{Pm} \sim \gamma(1, 0.1)$$

$$B_{Pm} \sim \gamma(1, 0.1)$$

$$A_{Pt} \sim \gamma(1, 0.1)$$

$$B_{Pt} \sim \gamma(1, 0.1)$$

$$M(k) \sim \gamma(1, 0.1)$$

$$\sigma(k) \sim \gamma(1, 0.1)$$

$$S_k = \frac{M(k)^2}{\sigma(k)^2}$$

$$R_k = \frac{M(k)}{\sigma(k)^2}$$

The γ distribution for the group level parameters is parameterised by its shape (S) and rate (R). Oberauer and colleagues (2017) placed priors on the mean and standard deviation of k (M_k, σ_k). Furthermore, the researchers proposed that the means and variances of the β distribution of the parameters is calculated as:

$$(9)\ M(Pm) = \frac{A_{Pm}}{A_{Pm}+B_{Pm}}$$

$$Var(Pm) = \frac{A_{Pm} \cdot B_{Pm}}{[A_{Pm} + B_{Pm}]^2[A_{Pm} + B_{Pm} + 1]}$$

$$M(Pt) = \frac{A_{Pt}}{A_{Pt} + B_{Pt}}$$

$$Var(Pt) = \frac{A_{Pt} \cdot B_{Pt}}{[A_{Pt} + B_{Pt}]^2[A_{Pt} + B_{Pt} + 1]}$$

In other words, the hierarchical Bayesian three-parameter mixture model (Oberauer et al., 2017) parametrises k, P_m, P_{guess}, and P_{nt} such that both individual differences and the group level memory parameters are accounted for under individual-specific and group-specific distributions. Arguably this allows for researchers to understand whether the group level memory parameters reflect individual feature retrieval. Thus, the group level parameters are reflected by the participants' individual memory parameters. Through a Bayesian approach, the hierarchical three-parameter mixture model allows for researchers to specify the distribution of the parameters on each level of its hierarchy. This highlights the flexibility of Bayesian hierarchical models, relative to the two- and three-parameter mixture models (Bays et al., 2009; Zhang & Luck, 2008). Additionally, fitting Oberauer and colleagues' (2017) model to the retrieved feature values in a continuous-report paradigm allows one to draw on Markov-Chain Monte-Carlo (MCMC) algorithms for estimating the posterior distribution of the memory parameter values. This allows one to determine the convergence of the model fitted relative to the data.

Moreover, fitting Oberauer and colleagues' (2017) model allows for researchers to analyse whether the difference between two conditions (i.e., experimental and control) is credible by determining whether the low and high highest-density intervals (HDI) overlap. For example, while comparing the experimental condition with the control (baseline) condition, if both the low HDI and high HDI are less than 0, or they are both greater than 0, then this is considered a credible effect. This is particularly advantageous when comparing the mean memory parameters (k, P_m, P_{guess}, and P_{nt}) between two conditions, as this offers researchers an insight as to whether the manipulation within the experimental condition impacts feature retrieval, relative to the control condition. Therefore, to understand the processing of single features and feature bindings, all the work in this thesis fit a hierarchical Bayesian three-parameter mixture model (Oberauer et al., 2017).

Appendix B
Chapter 2 Power Analysis

The power analyses determining the sample size for Experiment 1 and Experiment 2 of Chapter 2 are reported below. The power analysis R script is readily available on the Open Science Framework (OSF; https://osf.io/jr3eh/). The description of the power analyses reported below is reported in Stage 1 review of the registered report (accepted 1 February 2021) and is available on the OSF (https://osf.io/yrm83). For brevity the power analysis was not included in Stage 2 review, nor was it included in the published version of the registered report (https://doi.org/10.3758/s13414-023-02655-y).

Experiment 1

We determined the sample size of 24 participants, and the number of trials per condition by simulating 150 experiments based on the parameter estimates derived from fitting a hierarchical Bayesian three-parameter mixture model to the raw data of Souza and Oberauer (2017, Experiment 1A) whose design we closely follow in the current experiment. The analysis scripts and results of the simulations for both Experiments 1 and 2 can be found on the OSF. The rationale of this "power analysis" was to determine an appropriate number of participants and trials per condition to observe the predicted effects in at least 80% of the convergent simulations.

As explained previously, in Experiment 1 we predict that the central (but not peripheral) attention demand should increase extrinsic binding errors relative to the no-disruption baseline. We consider an effect as credible when the 95% highest density interval (HDI) of the comparison of interest (e.g., central versus no disruption) does not overlap with 0. The HDI gives a sense of how certain it is that the true effect lies in its range and its overlap with the null, thereby allowing us to draw inferences for our predictions. When fitting the three-parameter mixture model to Souza and Oberauer's raw data, we observed that the peripheral attention demand indeed had no impact on binding errors (mean estimate = 0.00, HDI range = -0.11 – 0.11). Conversely, although the central attention demand credibly impaired the probability that the target was recalled (mean estimate = 0.15, HDI range = 0.05 – 0.25), the effect on binding errors was not credible (mean estimate = -0.07, HDI range = -0.18 – 0.03). Given that other similar work has shown a credible effect of a central attention demand on binding errors (e.g., Zokaei et al., 2014), we wondered whether the demand was not strong enough to elicit an effect on binding errors. Indeed, there was no credible effect of a "low" central attention demand in Souza and Oberauer's Experiment 1B on the probability of recalling the target (mean estimate = 0.07, HDI range = -0.05 – 0.20) or the probability of binding errors (mean estimate = -0.08, HDI range = -0.21 – 0.06). Thus, in the current experiment we plan to adjust the pitch of the tones to make them more challenging to discriminate to increase the demand of the task (Portrat et al., 2008). Accordingly, we conducted the simulations with a mean difference of 0.15 in binding errors between the central attention and no disruption conditions. The results showed that 83% of the HDIs for the central attention disruption effect did not overlap with 0 (range *M* = 0.21, *SD* = 0.01), whereas 99% of the HDIs for the peripheral attention disruption effect overlapped with 0 (range *M* = 0.18, *SD* = 0.01).

Experiment 2

We will collect data from 24 participants, with the same criteria as in Experiment 1. We determined this sample size and number of trials per condition by using the same simulated experiments method from Experiment 1. Specifically, we conducted 150 simulated experiments using the parameter estimates derived from fitting a hierarchical Bayesian three-parameter mixture model to the younger adult data of Souza (2016), who showed that (extrinsic) retro-cues reduced binding errors relative to a no-cue baseline (mean estimate = -0.18, HDI range = -0.26 – -0.09). As in Experiment 1, the logic of this analysis was to check the percentage of convergent simulations that showed a credible retro-cue effect in binding errors with a given sample size and number of trials per condition. The results showed that 100% of the HDIs for the retro-cue effect in binding errors did not overlap with 0 (range M = 0.13, SD = 0.004), thus justifying the proposed sample size.

References

Adam, K. C. S., Vogel, E. K., & Awh, E. (2017). Clear evidence for item limits in visual working memory. *Cognitive Psychology*, *97*, 79–97. https://doi.org/10.1016/j.cogpsych.2017.07.001

Allen, R. J., Baddeley, A. D., & Hitch, G. J. (2006). Is the binding of visual features in working memory resource-demanding? *Journal of Experimental Psychology. General*, *135*(2), 298–313. https://doi.org/10.1037/0096-3445.135.2.298

Allen, R. J., Hitch, G. J., Mate, J., & Baddeley, A. D. (2012). Feature binding and attention in working memory: A resolution of previous contradictory findings. *Quaterly Journal of Experimental Psychology*, *65*(12), 2369–2383. https://doi.org/10.1080/17470218.2012.687384

Allen, R. J., Havelka, J., Falcon, T., Evans, S., & Darling, S. (2015). Modality specificity and integration in working memory: Insights from visuospatial bootstrapping. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*(3), 820–830. https://doi.org/10.1037/xlm0000058

Allen, R. J., Hitch, G. J., Mate, J., & Baddeley, A. D. (2012). Feature binding and attention in working memory: A resolution of previous contradictory findings. *The Quarterly Journal of Experimental Psychology*, *65*(12), 2369–2383. https://doi.org/10.1080/17470218.2012.687384

Arnicane, A., & Souza, A. S. (2021). Assessing the robustness of feature-based selection in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *47*(5), 731–758. https://doi.org/10.1037/xhp0000911

Astle, D. E., Summerfield, J., Griffin, I., & Nobre, A. C. (2012). Orienting attention to locations in mental representations. *Attention, Perception, & Psychophysics*, *74*(1), 146–162. https://doi.org/10.3758/s13414-011-0218-3

Atkinson, R. C., & Shiffrin, R. M. (1968). Human Memory: A Proposed System and its Control Processes. In K. W. Spence & J. T. Spence (Eds.), *Psychology of Learning and Motivation* (Vol. 2, pp. 89–195). Academic Press. https://doi.org/10.1016/S0079-7421(08)60422-3

Baddeley, A. (1996). Exploring the Central Executive. *The Quarterly Journal of Experimental Psychology Section A*, *49*(1), 5–28. https://doi.org/10.1080/713755608

Baddeley, A. D. (2000). The episodic buffer: A new component of working memory? *Trends in Cognitive Science*, *4*(11), 417–423. https://doi.org/10.1016/S1364-6613(00)01538-2

Baddeley, A. D. (2012). Working memory: Theories, models and controversies. *Annual Review Psychology*, *63*(1), 1–29. https://doi.org/10.1146/annurev-psych-12710-100422

Baddeley, A. D., Allen, R. J., & Hitch, G. J. (2011). Binding in visual working memory: The role of the episodic buffer. *Neuropsychologia*, *49*(6), 1393–1400. https://doi.org/10.1016/j.neuropsychologia.2010.12.042

Baddeley, A. D., & Hitch, G. J. (1974). Working Memory. In *Psychology of learning and motivation* (Vol. 8, pp. 47–89). Academic Press.

Baddeley, A., Hitch, G., & Allen, R. (2020). A Multicomponent Model of Working Memory. In R. Logie, V. Camos, & N. Cowan (Eds.), *Working Memory: The state of the science* (p. 0). Oxford University Press. https://doi.org/10.1093/oso/9780198842286.003.0002

Baddeley, A., Lewis, V., & Vallar, G. (1984). Exploring the Articulatory Loop. *The Quarterly Journal of Experimental Psychology Section A*, *36*(2), 233–252. https://doi.org/10.1080/14640748408402157

Barrouillet, P. N., & Camos, V. (2010). Working memory and executive control: A time-based resource sharing account. *Psychologica Belgica*, *50*(3/4), 353–382.

Bays, P. M. (2016). Evaluating and excluding swap errors in analogue tests of working memory. *Scientific Reports*, *6*(1), Article 1. https://doi.org/10.1038/srep19203

Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), 7–7. https://doi.org/10.1167/9.10.7

Bays, P. M., & Husain, M. (2008). *Dynamic Shifts of Limited Working Memory Resources in Human Vision*. https://doi.org/10.1126/science.1158023

Bays, P. M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia*, *49*(6), 1622–1631. https://doi.org/10.1016/j.neuropsychologia.2010.12.023

Beschin, N., Denis, M., Logie, R. H., & Della Sala, S. (2005). Dissociating mental transformations and visuo-spatial storage in working memory: Evidence from representational neglect. *Memory*, *13*(3–4), 430–434. https://doi.org/10.1080/09658210344000431

Brady, T. F., Störmer, V. S., & Alvarez, G. A. (2016). Working memory is not fixed-capacity: More active storage capacity for real-world objects than for simple stimuli. *Proceedings of the National Academy of Sciences*, *113*(27), 7459–7464. https://doi.org/10.1073/pnas.1520027113

Brissenden, J. A., & Somers, D. C. (2019). Cortico-Cerebellar Networks for Visual Attention and Working Memory. *Current Opinion in Psychology*, *29*, 239–247. https://doi.org/10.1016/j.copsyc.2019.05.003

Brown, L. A., & Brockmole, J. R. (2010). The role of attention in binding visual features in working memory: Evidence from cognitive ageing. *The Quarterly Journal of Experimental Psychology*, *63*(10), 2067–2079. https://doi.org/10.1080/17470211003721675

Camos, V., & Barrouillet, P. (2014). Attentional and non-attentional systems in the maintenance of verbal information in working memory: The executive and phonological loops. *Frontiers in Human Neuroscience*, *8*. https://www.frontiersin.org/articles/10.3389/fnhum.2014.00900

Camos, V., Johnson, M. R., Loaiza, V. M., Portrat, S., Souza, A. S., & Vergauwe, E. (2018). What is attentionl refreshing in working memory? *Annals of the New York Academy of Sciences*, *1424*(1), 19-32.

Carriedo, N., Corral, A., Montoro, P. R., Herrero, L., & Rucián, M. (2016). Development of the updating executive function: From 7-year-olds to young adults. *Developmental Psychology*, *52*, 666–678. https://doi.org/10.1037/dev0000091

Cocchini, G., Logie, R. H., Sala, S. D., MacPherson, S. E., & Baddeley, A. D. (2002). Concurrent performance of two memory tasks: Evidence for domain-specific working memory systems. *Memory & Cognition*, *30*(7), 1086–1095. https://doi.org/10.3758/BF03194326

Colom, R., Flores-Mendoza, C., Quiroga, M. Á., & Privado, J. (2005). Working memory and general intelligence: The role of short-term storage. *Personality and Individual Differences*, *39*(5), 1005–1014. https://doi.org/10.1016/j.paid.2005.03.020

Cowan, N. (1993). Activation, attention, and short-term memory. *Memory & Cognition*, *21*(2), 162–167. https://doi.org/10.3758/BF03202728

Cowan, N. (1998). *Attention and Memory: An Integrated Framework*. Oxford University Press.

Cowan, N. (1999). An Embedded-Processes Model of Working Memory. In A. Miyake & P. Shah (Eds.), *Models of Working Memory* (1st ed., pp. 62–101). Cambridge University Press. https://doi.org/10.1017/CBO9781139174909.006

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*(1), 87–185. https://doi.org/10.1017/s0140525x01003922

Cowan, N. (2016). *Working Memory Capacity: Classic Edition*. Routledge. https://doi.org/10.4324/9781315625560

Cowan, N., Blume, C. L., & Saults, J. S. (2013). Attention to attributes and objects in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(3), 731–747. https://doi.org/10.1037/a0029687

Cowan, N., & Morey, C. C. (2007). How Can Dual-Task Working Memory Retention Limits Be Investigated? *Psychological Science*, *18*(8), 686–688. https://doi.org/10.1111/j.1467-9280.2007.01960.x

Cowan, N., Morey, C. C., & Naveh-Benjamin, M. (2020). An Embedded-Processes Approach to Working Memory: How Is It Distinct From Other Approaches, and to What Ends? In R. Logie, V. Camos,

& N. Cowan (Eds.), *Working Memory: The state of the science* (p. 0). Oxford University Press. https://doi.org/10.1093/oso/9780198842286.003.0003

Darling, S., Allen, R. J., & Havelka, J. (2017). Visuospatial Bootstrapping: When Visuospatial and Verbal Memory Work Together. *Current Directions in Psychological Science*, *26*(1), 3–9. https://doi.org/10.1177/0963721416665342

Ecker, U. K. H., Lewandowsky, S., & Oberauer, K. (2014). Removal of information from working memory: A specific updating process. *Journal of Memory and Language*, *74*, 77–90. https://doi.org/10.1016/j.jml.2013.09.003

Ecker, U. K. H., Maybery, M., & Zimmer, H. D. (2013). Binding of extrinsic and intrinsic features in working memory. *Journal of Experimental Psychology: General*, *142*(1), 218–234. https://doi.org/10.1037/a0028732

Ecker, U. K. H., Oberauer, K., & Lewandowsky, S. (2014). Working memory updating involves item-specific removal. *Journal of Memory and Language*, *74*, 1–15. https://doi.org/10.1016/j.jml.2014.03.006

Elliott, E. M., Bell, R., Gorin, S., Robinson, N., & Marsh, J. E. (2022). Auditory distraction can be studied online! A direct comparison between in-Person and online experimentation. *Journal of Cognitive Psychology*, *34*(3), 307–324. https://doi.org/10.1080/20445911.2021.2021924

Elsley, J. V., & Parmentier, F. B. R. (2009). Is verbal–spatial binding in working memory impaired by a concurrent memory load? *The Quarterly Journal of Experimental Psychology*, *62*(9), 1696–1705. https://doi.org/10.1080/17470210902811231

Erdfelder, E., Auer, T.-S., Hilbig, B. E., Aßfalg, A., Moshagen, M., & Nadarevic, L. (2009). Multinomial Processing Tree Models. *Zeitschrift Für Psychologie / Journal of Psychology*, *217*(3), 108–124. https://doi.org/10.1027/0044-3409.217.3.108

Foungie, D., Cormeia, S. M., & Alvarez, G. A. (2013). Object-based benefits without object-based representations. *Journal of Experimental Psychology: General*, *142*(3), 621–626. https://doi.org/10.1037/a0030300

Gelman, A., Hwang, J., & Vehtari, A. (2014). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, *24*(6), 997–1016. https://doi.org/10.1007/s11222-013-9416-2

Griffin, I. C., & Nobre, A. C. (2003). Orienting Attention to Locations in Internal Representations. *Journal of Cognitive Neuroscience*, *15*(8), 1176–1194. https://doi.org/10.1162/089892903322598139

Gunseli, E., van Moorselaar, D., Meeter, M., & Olivers, C. N. L. (2015). The reliability of retro-cues determines the fate of noncued visual working memory representations. *Psychonomic Bulletin and Review*, *22*(1), 1334–1341. https://doi.org/10.3758/s13423-014-0796-x

Hajonides, J. E., van Ede, F., Stokes, M. G., & Nobre, A. C. (2020). Comparing the prioritization of items and feature-dimensions in visual working memory. *Journal of Vision*, *20*(8), 1–12. https://doi.org/10.1167/jov.20.8.25

He, Y., & Evans, A. (2010). Graph theoretical modeling of brain connectivity. *Current Opinion in Neurology*, *23*(4), 341. https://doi.org/10.1097/WCO.0b013e32833aa567

Henninger, F., Shevchenko, Y., Mertens, U. K., Kieslich, P. J., & Hilbig, B. E. (2021). lab.js: A free, open, online study builder. *Behavior Research Methods*. https://doi.org/10.3758/s13428-019-01283-5

Henninger, F., Shevchenko, Y., Mertens, U., Kieslich, P. J., & Hilbig, B. E. (2022). Lab.js: A free, open, online studybuilder. *Behavior Research Methods*, *54*, 556–573. https://doi.org/10.3758/s13428-019-01283-5

Heuer, A., & Schubö, A. (2016). The Focus of Attention in Visual Working Memory: Protection of Focused Representations and Its Individual Variation. *PLOS ONE*, *11*(4), e0154228. https://doi.org/10.1371/journal.pone.0154228

Hitch, G. J., Allen, R. J., & Baddeley, A. D. (2020). Attention and binding in visual working memory: Two forms of attention and two kinds of buffer storage. *Attention, Perception, & Psychophysics*, *82*(1), 280–293. https://doi.org/10.3758/s13414-019-01837-x

Johnson, J. C., & Pashler, H. (1990). Close binding of identity and location in visual feature perception. *Journal of Experimental Psychology: Human Perception and Performance*, *16*(4), 843–856. https://doi.org/10.1037//0096-1523.16.4.843

Kessler, Y., & Meiran, N. (2008). Two dissociable updating processes in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(6), 1339–1348. https://doi.org/10.1037/a0013078

Kessler, Y., Rac-Lubashevsky, R., Lichtstein, C., Markus, H., Simchon, A., & Moscovitch, M. (2015). Updating visual working memory in the change detection paradigm. *Journal of Vision*, *15*(9), 1–12. https://doi.org/10.1167/15.9.18

Kim, H., Smolker, H. R., Smith, L. L., Banich, M. T., & Lewis-Peacock, J. A. (2020). Changes to information in working memory depend on distinct removal operations. *Nature Comminications*, *11*(1), 1–14. https://doi.org/10.1038/s41467-020-2008

Kirmsse, A., Zimmer, H. D., & Ecker, U. K. H. (2018). Age-related changes in working memory: Age affects relational but not conjunctive feature binding. *Psychology and Aging*, *33*(3), 512–526. https://doi.org/10.1037/pag0000249

Lange, K., Kühn, S., & Filevich, E. (2015). "Just Another Tool for Online Studies" (JATOS): An Easy Solution for Setup and Management of Web Servers Supporting Online Studies. *PLOS ONE*, *10*(6), e0130834. https://doi.org/10.1371/journal.pone.0130834

Langerock, N., Vergauwe, E., & Barrouillet, P. (2014). The maintenance of cross-domain associations in the episodic buffer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*(4), 1096–1109. https://doi.org/10.1037/a0035783

Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2012). Neural Evidence for a Distinction between Short-term Memory and the Focus of Attention. *Journal of Cognitive Neuroscience*, *24*(1), 61–79. https://doi.org/10.1162/jocn_a_00140

Lewis-Peacock, J. A., & Postle, B. R. (2012). Decoding the internal focus of attention. *Neuropsychologia*, *50*, 470–478. https://doi.org/10.1016/j.neuropsychologia.2011.11.006

Lewis-Peacock, J. L., Kessler, Y., & Oberauer, K. (2018). The removal of information from working memory. *Annals of the New York Academy of Sciences*, *1424*(1), 33–44. https://doi.org/10.1111/nyas.13714

Li, A. Y., Liang, J. C., Lee, A. C. H., & Barense, M. D. (2020). The validated circular shape space: Quantifying the visual similarity of shape. *Journal of Experimental Psychology. General*, *149*(5), 949–966. https://doi.org/10.1037/xge0000693

Li, Q., & Saiki, J. (2015). Different effects of color-based and location-based selection on visual working memory. *Attention, Perception, & Psychophysics*, *77*(2), 450–463. https://doi.org/10.3758/s13414-014-0775-3

Loaiza, V. M., Cheung, H. W., & Goldenhaus-Manning, D. T. (2023). What you don't know can't hurt you: Retro-cues benefit working memory regardless of prior knowledge in long-term memory. *Psychonomic Bulletin & Review*. https://doi.org/10.3758/s13423-023-02408-w

Loaiza, V. M., & Souza, A. S. (2019a). An age-related deficit in preserving the benefits of attention in working memory. *Psychology and Aging*, *34*(2), 282–293. https://doi.org/10.1037/pag0000326

Logie, R. H. (1995). *Visuo-spatial Working Memory*. Psychology Press.

Logie, R. H. (2003). Spatial and Visual Working Memory: A Mental Workspace. In *Psychology of Learning and Motivation* (Vol. 42, pp. 37–78). Academic Press. https://doi.org/10.1016/S0079-7421(03)01002-8

Logie, R. H., & Pearson, D. G. (1997). The Inner Eye and the Inner Scribe of Visuo-spatial Working Memory: Evidence from Developmental Fractionation. *European Journal of Cognitive Psychology*, *9*(3), 241–257.

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279–281. https://doi.org/10.1038/36846

Luck, S. J., & Vogel, E. K. (2013). Visual working memory capacity: From psychophysics and neurobiology to individual differences. *Trends in Cognitive Sciences*, *17*(8), 391–400. https://doi.org/10.1016/j.tics.2013.06.006

Luria, R., & Vogel, E. K. (2011). Shape and color conjunction stimuli are represented as bound objects in visual working memory. *Neuropsychologia*, *49*(6), 1632–1639. https://doi.org/10.1016/j.neuropsychologia.2010.11.031

Morey, C. C., & Bieler, M. (2013). Visual short-term memory always requires general attention. *Psychonomic Bulletin & Review*, *20*(1), 163–170. https://doi.org/10.3758/s13423-012-0313-z

Morey, R. D., Rouder, J. N., Jamil, T., Urbanek, S., Forner, K., & Ly, A. (2022). *BayesFactor: Computation of Bayes Factors for Common Designs* (0.9.12-4.4) [Computer software]. https://cran.r-project.org/web/packages/BayesFactor/index.html

Oberauer, K. (2002). Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *28*(3), 411–421. https://doi.org/10.1037//0278-7393.28.3.411

Oberauer, K. (2019). Working memory and attention—A conceptual analysis and review. *Journal of Cognition*, *2*(1), 1–23. https://doi.org/10.5334/joc.58

Oberauer, K., Stoneking, C., Wabersich, D., & Lin, H.-Y. (2017). Hierarchical Bayesian measurement models for continuous reproduction of visual features from working memory. *Journal of Vision*, *17*(5), 11–38. https://doi.org/10.1167/17.5.11

Palladino, P., Cornoldi, C., De Beni, R., & Pazzaglia, F. (2001). Working memory and updating processes in reading comprehension. *Memory & Cognition*, *29*(2), 344–354. https://doi.org/10.3758/BF03194929

Pertzov, Y., & Husain, M. (2014). The privileged role of location in visual working memory. *Attention, Perception & Psychophysics*, *76*(7), 1914–1924. https://doi.org/10.3758/s13414-013-0541-y

Plummer, M. (2016). *rjags: Bayesian graphical models using MCMC* [Computer software]. https://cran.r-project.org/web/packages/rjags/rjags.pdf

Plummer, M., Stukalov, A., & Denwood, M. (2022). *rjags: Bayesian Graphical Models using MCMC* (Version 4-13) [Computer software]. https://CRAN.R-project.org/package=rjags

Portrat, S., Barrouillet, P., & Camos, V. (2008). Time-related decay or interference-based forgetting in working memory? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(6), 1561–1564. https://doi.org/10.1037/a0013356

Rajsic, J., & Wilson, D. E. (2014). Asymmetrical access to color and location in visual working memory. *Attention, Perception, & Psychophysics*, *76*(7), 1902–1913. https://doi.org/10.3758/s13414-014-0723-2

Ricker, T. J., & Vergauwe, E. (2022). Boundary conditions for observing cognitive load effects in visual working memory. *Memory & Cognition*. https://doi.org/10.3758/s13421-022-01320-3

Schneegans, S., & Bays, P. M. (2017). Neural Architecture for Feature Binding in Visual Working Memory. *Journal of Neuroscience*, *37*(14), 3913–3925. https://doi.org/10.1523/JNEUROSCI.3493-16.2017

Schneegans, S., & Bays, P. M. (2019). New perspectives on binding in visual working memory. *British Journal of Psychology*, *110*(2), 207–244. https://doi.org/10.1111/bjop.12345

Shepherdson, P., Hell, L., & Oberauer, K. (2022). How does visual working memory solve the binding problem? *Journal of Experimental Psychology: Human Perception and Performance*, *48*, 1137–1152. https://doi.org/10.1037/xhp0001044

Souza, A. S. (2016). No age deficits in the ability to use attention to improve visual working memory. *Psychology and Aging*, *31*(5), 456–470. https://doi.org/10.1037/pag0000107

Souza, A. S., & Oberauer, K. (2016). In search of the focus of attention in working memory: 13 years of the retro-cue effect. *Attention, Perception, & Psychophysics*, *78*(7), 1839–1860. https://doi.org/10.3758/s13414-016-1108-5

St Clair-Thompson, H., Stevens, R., Hunt, A., & Bolder, E. (2010). Improving children’s working memory and classroom performance. *Educational Psychology*, *30*(2), 203–219. https://doi.org/10.1080/01443410903509259

Swan, G., & Wyble, B. (2014). The binding pool: A model of shared neural resources for distinct items in visual working memory. *Attention, Perception, & Psychophysics*, *76*(7), 2136–2157. https://doi.org/10.3758/s13414-014-0633-3

Teng, C., & Kravitz, D. J. (2019). Visual working memory directly alters perception. *Nature Human Behaviour*, *3*(8), 827–836. https://doi.org/10.1038/s41562-019-0640-4

Treisman, A. (1977). Focused attention in the perception and retrieval of multidimensional stimuli. *Perception & Psychophysics*, *22*(1), 1–11. https://doi.org/10.3758/BF03206074

Treisman, A. (1999). Solutions to the Binding Problem: Progress through Controversy and Convergence. *Neuron*, *24*(1), 105–125. https://doi.org/10.1016/S0896-6273(00)80826-0

Treisman, A. M., & Gelade, G. (1980). A Feature Intergration Theory of Attention. *Cognitive Psychology*, *12*(1), 97–136. https://doi.org/10.1016/0010-0285(80)90005-5

Ueno, T., Allen, R. J., Baddeley, A. D., Hitch, G. J., & Saito, S. (2011). Disruption of visual feature binding in working memory. *Memory & Cognition*, *39*(1), 12–23. https://doi.org/10.3758/s13421-010-0013-8

van Geldorp, B., Parra, M. A., & Kessels, R. P. C. (2015). Cognitive and neuropsychological underpinnings of relational and conjunctive working memory binding across age. *Memory*, *23*(8), 1112–1122. https://doi.org/10.1080/09658211.2014.953959

Vergauwe, E., Barrouillet, P., & Camos, V. (2009). Visual and spatial working memory are not that dissociated after all: A time-based resource-sharing account. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*(4), 1012–1028. https://doi.org/10.1037/a0015859

Vergauwe, E., Langerock, N., & Barrouillet, P. (2014). Maintaining information in visual working memory: Memory for bindings and memory for features are equally disrupted by increased attentional demands. *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale*, *68*(3), 158–162. https://doi.org/10.1037/cep0000025

Vogel, E. K., Woodman, G. F., & Luck, S. J. (2001). Storage of Features, Conjunctions, and Objects in Visual Working Memory. *Journal of Experimental Psychology: Human Perception and Performance*, *27*(1), 92–114. https://doi.org/10.1037/0096-1523.27.1.92

Watanabe, S. (2021). WAIC and WBIC for mixture models. *Behaviormetrika*, *48*(1), 5–21. https://doi.org/10.1007/s41237-021-00133-z

Wheeler, M. E., & Treisman, A. M. (2002). Binding in short-term memory. *Journal of Experimental Psychology: General*, *131*(1), 48–64. https://doi.org/10.1037//0096-3445.131.1.48

Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, *4*(12), 1120–1135. https://doi.org/10.1167/4.12.11

Wolfe, J. M. (2020). Forty years after feature integration theory: An introduction to the special issue in honor of the contributions of Anne Treisman. *Attention, Perception, & Psychophysics*, *82*(1), 1–6. https://doi.org/10.3758/s13414-019-01966-3

World Medical Association. (2013). World Medical Association Declaration of Helsinki: Ethical Principles for Medical Research Involving Human Subjects. *JAMA*, *310*(20), 2191–2194. https://doi.org/10.1001/jama.2013.281053

Ye, C., Hu, Z., Ristaniemi, T., Gendron, M., & Liu, Q. (2016). Retro-dimension-cue benefit in visual working memory. *Scientific Reports*, *6*(1), Article 1. https://doi.org/10.1038/srep35573

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), Article 7192. https://doi.org/10.1038/nature06860

Zokaei, N., Heider, M., & Husain, M. (2014). Attention is Required for Maintenance of Feature Binding in Visual Working Memory. *Quarterly Journal of Experimental Psychology*, *67*(6), 1191–1213. https://doi.org/10.1080/17470218.2013.852232