An Exploration of the Role of Temporal Information in Hierarchical Working Memory Representations and their Influences on Visual Attention

Mr Craig Peter Adams Arnold

Department of Psychology

Royal Holloway, University of London

Submitted for the Degree of Doctor of Philosophy, September 2020

Declaration

I, Craig Arnold, hereby declare that this thesis and the work presented in it is entirely my own. Where I have consulted the work of others, this is always clearly stated.

Signed: _____Craig Arnold_________________

Date: ____15/09/2020____________________

ACKNOWLEDGEMENTS

First and foremost, I need to thank my wonderful supervisor Prof. Polly Dalton for four years of constant support and encouragement. She never gave up, even when I had, and there is no way I would have completed this without her mentorship.

I also want to thank my partner Edo who has been with me for three of these four years and had to endure many cancelled plans and the near endless repetition of the phrase “after the PhD”. I would like to thank him for his support throughout, and just being there when I needed him most.

I also need to thank the wider group of people who have helped me get through these four years. From RHUL I would like to thank my advisor Steve for his guidance and input, and Sandra and Stuart for their encouragement and advice early on. I also need to thank the wider gang of PhD students for providing such an excellent source of support and friendship. A special thanks needs to go out to David, for the many discussions into the late evening and Rachel L for being there to listen when I needed to vent.

Outside of the PhD itself, I also need to thank Victoria for helping me to ‘survive’ the first year, and Egham, as without her there probably would not have been subsequent years. I would also not have been able to cope with those subsequent years if I did not have such great people in my life to balance things out. A special thanks must go to Gian, Anne-Cecilie and Astrid for making Reading more enjoyable than it really should have been.

Although this PhD has been only four years, the journey to this point goes back much farther, and it would be a disservice not also to acknowledge those who helped get me to the PhD in the first place. Most obviously I would like to thank both of my parents Peter and Sheila for their constant support, encouraging me to strive further, and for ensuring I had as many options open to me as possible (as well as coming to visit the far-off places I kept moving to). Also, a huge thanks to my sister Jo and brother-in-law Jimmy for helping keep me grounded and not take things so seriously.

I have been fortunate to have had a series of amazing academic mentors, each of whom shaped me and my ways of thinking in different ways. Primarily these have included Mrs Aline Marshall, Prof Verity Brown and Dr David Tait, Dr Bob Kentridge and Dr Cristiana Cavina-Pratesi, Prof Kia Nobre and Prof Carlo Marzi, and Dr Christina Howard and Dr Matthew Belmonte. In addition to Prof Dalton, these people made me into the scientist I am

today and, even though I am now pursuing alternative career paths, their influence and teachings remain as relevant as ever.

There was a period that I had left academia and all but given up on the idea of applying for PhDs. However, I decided to apply to one more Research Assistantship, with Dr Howard and Dr Belmonte, and it is only thanks to my experience there that my passion for research was reignited. Beyond my academic supervisors, this was also largely down to the friends I made in the CMDA lab. Josh, Jammy, Rob, Sarah and Mark, as well as others – thank you.

Although it may have been the people in Nottingham that rekindled the passion for research, it was in Durham that that passion was first ignited. In addition to Dr Kentridge and Dr Cavina Pratesi, I have to thank the MSc Cog Neuro et al. group. I remain forever grateful at how lucky we were to meet each other, forging friendships that continue to this day. Thank you to Carys, Misha, Lindsey, Pete, Alice, Kay, Cath, Alex, Cat, Lizzy and Eliza. A special thanks needs to go to Misha, Carys and Lindsey for pushing me all these years not to give up, especially during some of the more difficult periods. Thank you also to Nick and George in Oxford for acting as sounding boards to my early rough ideas and theories, and for helping me to secure my first publication.

I need to give a very special thanks also to Jamie Griggs. Without his support during the final months of my undergraduate degree, I would not have got the grades needed to pursue further study nor an academic career. I owe all of this to those weeks in his room, where he took time out of his own revision to help me catch up. Without his generosity, none of this would have been possible. It remains a major loss to science that he never got the opportunity to make his mark. He was a brilliant thinker, and wonderful friend, and his memory has continued to inspire me to work harder and makes me grateful for the opportunities I have.

This PhD, and the road leading up to it, is impossible to separate from other aspects of life. In addition to all the people above who have had some direct role in my academic and professional journey, there are all the other people along the way supporting me in other aspects of my life, making the tough times manageable, and helping to remind me what's really important. Some specific thanks need to go to my extended family, especially my grandparents, Aunt Pauline, and cousins, for reminding me to be grateful for all I have, and showing me how not to take life so seriously. A specific thank you also to Tom for both putting up with, and supporting me for four years of pursuit of a career in academia. Other thanks need to go to Tara, Shane, Ian, Matthew, Dan, James, Andy, Casey, Katie, Lucy, and countless others. Even though often we may have lost touch, I remain thankful for the support and influence each offered me at various stages of my life, all of which have ultimately shaped who I am today, including (however indirectly) completing this PhD.

To anyone not listed here, it is not that you have been forgotten. I would happily list each and every person who has supported me in so many various ways, and offer thanks to each, but it’s not really feasible here.

I suppose in summary, to all of you, I would like to say thank you.

Abstract

This thesis investigates how visual working memory effects may differ at different levels of a hierarchically structured system. Traditional visual working memory research has focussed on memory for exact features of individual objects. However, working memory can also include other types of information, such as the 'average' feature value across a group of stimuli, or the relative spatial relationships between stimuli. Most studies have used static displays of stimuli presented simultaneously, so this thesis aimed to investigate whether similar mechanisms could operate across stimuli presented sequentially, though this thesis failed to replicate previous results even when presenting simultaneous stimuli. It also investigated whether group level representations are observed using tactile stimuli, based on evidence of integration across tactile inputs, but fails to replicate these effects.

A second research topic was that of memory guided attention. Holding an item in memory can bias the allocation of attention in favour of new incoming information that shares similar features. However, it remained unclear how the degree of similarity between the new items and memory item modulated such guidance. This thesis attempted to replicate such guidance effects using orientation as the critical feature, to allow similarity to be quantified better. However, no memory guidance effects were observed under these conditions. Finally, so far previous studies had also primarily used static arrays of simultaneous stimuli, so this thesis investigated whether such guidance effects could affect the allocation of temporal attention, though again no such evidence was observed.

The thesis furthers our understanding of the limits of the interplay between attention and memory across time, by failing to find evidence of previously established mechanisms under new conditions, including novel feature dimensions, sequential presentation, and tactile stimuli and thus questioning the generalisability of the previous research in this area.

Table of Contents

List of Figures

List of Tables

General Introduction

The overarching theme of this thesis is the interplay between attention and working memory. These terms will be described in more detail below but, broadly speaking, attention can be thought of as the mechanism that allows one to focus on specific information, and working memory can be defined as the mechanism that allows for some perceptual information to be maintained even after it is no longer physically present. It is well understood that attention plays a key role in the formation and maintenance of working memory representations, and that these can, in turn, influence the deployment of attention across new information. The key theme this thesis aims to investigate is how this interplay operates at different levels (or hierarchies) of representation. Different levels of representation can be thought of as different levels of specificity within a scene. For example, even if one cannot recall a specific feature of an object, one might still remember ‘something’ else about it, or even if one cannot recall specific information about every item in a group, one might still remember information about the group as a whole. Most of the current work focuses on the visual modality but Chapter 3 also includes some work investigating similar questions in the tactile domain.

This first chapter will outline relevant previous research and the broad rationale for the general questions considered, whereas the introductions to each experimental chapter will include more focussed discussions relevant to those studies specifically. The first issue that will be addressed is to define how the terms “attention” and “working memory” will be used throughout the thesis, as there is often confusion within the literature from multiple competing definitions of these terms. However, it should be emphasised early on that these two mechanisms are strongly related, and the rest of this general introduction will then summarise the existing literature on this interplay in more detail and outline the key questions that the thesis aims to answer.

Attention

Attention is broadly considered to be the mechanism by which we are able to focus on specific inputs while ignoring others (Posner, 1980). A common framework for thinking about this ‘focus’ of attention, especially visual attention, has been as a spotlight (analogous to those used to highlight specific areas of the stage in theatre; Posner, Snyder, & Davidson, 1980) that can be moved around the scene. This approach typically considered that whatever this spotlight is aimed at is ‘attended’ and anything outside of it is ‘unattended’. The focus of this ‘beam’ need not be a fixed size though and may become more focused or diffuse, depending on task demands (“zoom-lens model”; Eriksen & St James, 1986). There are important problems with this analogy however, as will be outlined later, and more recently the ‘biased competition model’ of attention (Desimone & Duncan, 1995) has become a more dominant approach. Its core concept is that attention may not be a property that something ‘has’ or ‘has not’, but instead attentional effects arise from the relative biasing of inputs at the expense of others when competing for limited representation resources.

A central concept of the biased competition explanation of attention is that different stages of processing of visual information have different ‘capacities’ (Hochstein & Ahissar, 2002). For example, representing individual ‘edges’, or combining these edges into shapes, or shapes into objects, or objects across groups, and so on. At any stage where inputs contain more information than can be represented by the next stage of a system, those inputs ‘compete’ and the ‘losing’ information does not progress to further levels of representation. Sometimes some of these inputs may simply be stronger due to intrinsic properties such as being more salient (Kerzel & Schönhammer, 2013; Yantis, 2005). However, we are also able to focus voluntarily on other aspects of the scene than just the most salient stimulus, such as less salient but more behaviourally relevant stimuli (Indovina & MacAluso, 2007), or even on specific locations in a blank area where we anticipate that an important stimulus may appear (Posner, 1980).

These different types of attentional orienting are typically categorized into two main types, endogenous and exogenous (Posner, 2016). Endogenous (or top-down) attention is when one voluntarily and purposefully shifts the focus of their attention from one input to another in line with their current goals (such as when trying to focus on what a friend is saying at a party while ignoring the loud background noise). Exogenous (or bottom-up) attention refers to the situation where attention is ‘captured’ involuntarily, such as by something sudden and unexpected, yet unrelated to one’s ongoing task goals (for example, when distracted from an interesting conversation at the party by a loud smash as someone drops their glass). However, this traditional dichotomy is now considered by many to be too restrictive and cannot account for other forms of biasing, such as towards stimuli that have previously been associated with reward but are not relevant to the current task (so are not intrinsically salient, but are not in line with current task goals either; Awh, Belopolsky, & Theeuwes, 2012).

At a neural level, this biasing is achieved via feed-back connections that allow ‘higher’ areas to enhance or inhibit the strength of competing inputs in order to bias some relative to others (Desimone & Duncan, 1995; Gilbert & Sigman, 2007; Miconi & VanRullen, 2016). Although this initially seems an intuitive mechanism, there are a couple of ambiguities and controversies that are still often confused or neglected in the wider literature, but as they are assumed in this thesis, it is important to note them here.

For example, the presence of competition at a given stage of processing does not automatically indicate the potential for direct attentional biasing at that stage. Although competition could occur at any stage in the neural processing pathway where there is a reduction in capacity from one stage to the next, attentional biasing may be limited by the neural architecture of these feedback connections. Such feedback may not exist (and therefore direct attentional biasing may not be possible) at every instance where competition occurs.

For example, in electroencephalography (EEG), one of the visually evoked event-related potentials (ERP), the C1, peaks before 100ms and is thought to reflect early feed-forward processing (Alilović, Timmermans, Reteig, Van Gaal, & Slagter, 2019). Many previous studies have claimed that this early ERP is not modulated by top-down attentional conditions (Alilović et al., 2019; Fu, Fedota, Greenwood, & Parasuraman, 2010), while slightly later vERPs such as the P1 (peaking approximately 110ms after stimulus onset) do seem to vary under different attentional conditions (Hillyard & Anllo-Vento, 1998). These authors have therefore argued that the neural mechanisms underlying some of the earliest processing of incoming information cannot be biased by top-down feedback connections (though note that there is evidence that even C1 may be subject to attentional modulation under some situations; Slotnick, 2018).

Similarly, there is debate about the extent to which these top-down attentional biases can modulate neuronal activity in early visual areas. For example, some previous studies have demonstrated that when two stimuli are presented simultaneously, activity of neurons in areas such as V2 and V4 differed depending on which stimulus was being attended when both stimuli were within the neuron’s receptive field (and so in competition) but not when only one was presented within the receptive field (so no biasing of competition was necessary; Luck, Chelazzi, Hillyard, & Desimone, 1997). Different attentional states did not appear to affect activity in V1 neurons, though it was noted that given the small size of the receptive fields, it was not possible to present two competing stimuli within the same receptive field. However, other approaches have argued that although attentional modulation of V1 may not always be reliably detected at an individual neuron level, spatial attention may still modulate overall activity in V1 at a population level (Sengpiel & Hübener, 1999). For example, focussing on one hemifield over the other may alter activity patterns across early visual regions that represent the different hemifields.

So far the description of attention has focussed on the biasing of spatially-separated inputs that are competing simultaneously for limited processing. However, attentional prioritisation has also been shown to operate between stimuli that are separated in time (temporal attention) as opposed to space. For example, when attempting to maintain focus on a task for a sustained period of time, one's attention (and associated efficiency of processing incoming information) naturally waxes and wanes over time (Macdonald, Mathan, & Yeung, 2011). In a similar manner as there are limits on how attention can be deployed across other feature dimensions such as space, there appear to be limitations on efforts to attend over time.

In addition to natural variation over time, attention can also be purposefully deployed to particular moments in time in response to task conditions. For example, if it is known that a task critical stimulus always appears one second after a cue stimulus, it would be most efficient to focus attention during that critical time period relative to other time periods where the stimulus is unlikely to appear. Similarly, if it is known when a distracting stimulus will be presented, it would improve task performance to inhibit information during that time period. These effects have been observed both in terms of improved behavioural performance (such as reaction time; Griffin, Miniussi, & Nobre, 2001) and using electrophysiological measures (such as occipital alpha; Rohenkohl & Nobre, 2011).

Although temporal attention can also interact with other forms of attention, such as spatial attention (Kizuk, Sayeed & Mathewson, 2017), temporal attention effects can be observed even when only one stimulus is visible at a time (Griffin et al., 2001). This may appear surprising, given that attention has been described as the relative biasing of competing inputs, and it may not be clear what competition could exist if only a single stimulus was present. However, attentional feed-back mechanisms do not require there to be actual inputs present and it is well known that biases in sensitivity can occur in anticipation of actual stimulation (Rohenkohl & Nobre, 2011). In some circumstances, as will be outlined in further detail in the working memory subsection, there may also be competition between new external information and existing internal recurrent processing, with attention switching between biasing or inhibiting new information in favour of existing representations (Chun, 2011). It is worth clarifying a common misunderstanding in the literature, that periods of 'no attention' or 'without attention' (e.g. Mack, Erol, Clarke, & Bert, 2016; Moore & Egeth, 1997; Moore, 2001) are often actually instances where biasing between competing inputs has been maximised (strongly biasing attention in favour of one input, to test what residual processing remains of negatively biased stimuli). Instead, a more parsimonious use of the term 'no attention' might not be during periods of large inhibition, but merely a lack of top-down biasing of competing inputs, leaving competition to be decided by other factors (this is the definition of 'no attention' assumed in this thesis). Therefore, the relative sensitivities (both positive and negative) of the visual system to various features, and even simply to external information at all, can vary over time, sometimes in a focussed and purposeful manner (attended/inhibited) and other times in an unguided and bottom-up manner (unattended).

Working Memory, or Visual Short Term Memory (VSTM)

Experiments in this thesis employ paradigms that include measuring participants' robust visual short-term memory, which will be referred to as 'working memory' throughout this thesis. It should be emphasized that the term 'working memory' has been applied broadly across the literature but with different underlying definitions and assumptions. In this thesis, the term 'working memory' is used more in relation to the Cowan (Cowan, 1998) approach to working memory, rather than the Baddeley (Baddeley & Hitch, 1974) model specifically. However, ultimately this thesis adopts a similar approach to other groups and defines working memory operationally as memory that can last for longer than one second and whose information is minimally disrupted by intervening masking stimuli (e.g. Sligte, Scholte, & Lamme, 2008; Soto, Mäntylä, & Silvanto, 2011; Vandenbroucke, Sligte, & Lamme, 2011).

Many researchers posit other characteristics, such as that the remembered information must be conscious in order for the process to qualify as working memory (Bor & Seth, 2012) or that the information must be manipulated in mind. However, most studies of working memory do not explicitly test for these criteria and it remains ambiguous whether such characteristics are truly critical requirements, rather than associated phenomena often observed due to the types of tasks commonly used. For example, there is evidence of working-like memory (persisting for longer than one second and surviving intervening masks) in the absence of reports of conscious awareness and it is debated whether such a phenomenon is evidence of unconscious working memory (Soto et al., 2011) or whether it must be a different type of memory (Astle, Nobre, & Scerif, 2010) given that many posit working memory must be conscious by definition (Bor & Seth, 2012). Similarly, though information encoded into the working memory 'workspace' (Baars & Franklin, 2003) can be mentally manipulated, it seems counterintuitive to classify any such representation as a memory only after it has been meaningfully altered from its original state, and not when it is at its most similar to the to-be-remembered information. Working memory may typically be accompanied by imagery, and indeed the two mechanisms appear to show a strong relation to each other (Tong, 2013). However, a more parsimonious account might be for such a representation to be classified as memory before any such manipulation (potentially with concurrent imagery) and as mental imagery after. Note that I do not object to a representation still being classified as working memory even after it has been manipulated, only to the claim that such a manipulation is necessary.

It should be emphasized that neither of the stated requirements, of persisting for longer than one second or through intervening masking stimuli, are prerequisite conditions for a working memory representation to form either. Working memory encoding likely starts shortly after the target stimulus onset and therefore a working memory representation may be formed and exist even before the stimulus has been physically removed (Moore, 2001; Tsubomi, Fukuda, Watanabe, & Vogel, 2013; Vogel, Woodman, & Luck, 2006; though some researchers may be uncomfortable with the semantics of labelling such a representation as 'memory'), and any such representation can occur with or without a subsequent masking stimulus. For clarity, this thesis adopts the theoretical approach that the distinction between memory and perception is not whether the stimulus is still present or not, but whether one is actively attending to external information and using it to integrate or overwrite an internal representation (in which case the process is considered to be perception) or instead whether one is attempting to inhibit or ignore external information in favour of maintaining internal recurrent processing of an existing representation (in which case the process is considered to be memory; though I recognise that this definition is somewhat controversial and not widely adopted). The position is therefore not that the stated criteria are the minimum required for such a representation to form, but instead are the criteria required to rule out effects from other forms of short-term memory (particularly iconic memory and fragile working memory), to be confident that any observed effects can be ascribed to working memory specifically.

A full description of these other types of memory is beyond the scope of the current review but a brief outline is important as they have major implications for the design of the paradigms used. Iconic memory is thought to be driven primarily by the slow decay of retinal photoreceptors, such that even once the physical stimuli have disappeared, the pattern of activation takes a short time to decay and during this period (200-500ms) selective attention can still access and encode this information (Sperling, 1960). This gives iconic memory a seemingly high capacity (it is often shown to operate well over arrays of even 12 alphanumeric stimuli) but its duration is very brief. Importantly, it is also easily disrupted and occurs mainly when the critical stimuli simply disappear and a uniform background remains. If another stimulus, or even just a flash of light (Sperling, 1960; Vandenbroucke et al., 2011), is presented following the offset of the memory stimulus, it appears that the active change in photoreceptor activity arising from the processing of these new stimuli erases the activity from the previous stimuli to an extent that such information can no longer be accessed. For this reason, working memory paradigms will typically use delay periods of greater than one second and will often include some form of intervening mask in order to ensure that any remaining memory effects cannot be influenced by any lingering iconic memory mechanisms.

An additional, more recent consideration is that of fragile working memory (Sligte, Wokke, Tesselaar, Steven Scholte, & Lamme, 2011; Vandenbroucke et al., 2011). Fragile working memory can persist for several seconds, and through an intervening flash of light, distinguishing it from iconic memory. However, it also appears to have a relatively high capacity (the amount of information that can be maintained at once), distinguishing it from traditional 'robust' working memory, which has a severely limited capacity of only 3-4 items-worth of information. Therefore, to isolate the effects of robust working memory only, it may be necessary to ensure the intervening mask stimulus has sufficiently confusing featural information, rather than a simple uniform flash. However, it should also be noted that studies demonstrating fragile working memory seem to require extensive training sessions beforehand (Sligte et al., 2011; Vandenbroucke et al., 2011) and the authors acknowledge that these effects are observed under artificial conditions (though they argue that the ability itself may exist naturally, only that participants may require practice and training to make use of it consistently enough for the effects to be observed with sufficient reliability). Many studies have not adopted such an approach (either using a simple flash mask, or using delay periods sufficiently long to rule out iconic memory; Luck & Vogel, 1997) and have still replicated the low capacity of robust working memory, with no apparent evidence of fragile working memory. Nevertheless, the studies in this thesis adopt a mask with confusable features, to reduce any such possible sources of ambiguity in interpreting observed effects.

There is much debate within the literature on whether the number of items that can be maintained in working memory is affected by the complexity of each item, such that fewer complex stimuli could be remembered by comparison with simple stimuli (Bays, Catalao, & Husain, 2009; Bays & Husain, 2008; Ma, Husain, & Bays, 2014), or whether the memory capacity reflects a limited number of ‘slots’ for individual items, largely independent of the complexity of each item (Awh, Barton, & Vogel, 2007; W. Zhang & Luck, 2008). It is possible that capacity is instead a combination of both slot-based mechanisms and continuous mechanisms. For example, it may be that for simple features, capacity operates in a slot-based manner, such that a few basic features can be grouped and recalled with little extra effort than a single simple feature, but if the stimuli get sufficiently complex that they cannot be grouped easily that additional effort to recall the extra details may either reduce the number of ‘slots’ available or else assign multiple slots to different aspects of the same item. This debate has little direct relevant for the current thesis however, and throughout the thesis any mention of working memory capacity in terms of items is intended as items-worth of information, and is neutral on whether this information if comprised of only distinct slots, of a continuous measure, or some combination of both.

Traditional working memory paradigms tended to use categorical stimuli such as alphanumeric characters (Ngiam, Khaw, Holcombe, & Goodbourn, 2018) and highly distinct colours (such as red, blue, green, yellow; Vogel et al., 2006), which allow responses to be classified as either correct or incorrect. However, internal representations are subject to noise, interference and uncertainty, meaning that this dichotomy of either a perfect match or completely forgotten is unlikely to match the reality of the underlying mechanisms. When using stimuli with clear and distinct categories though, even imperfect memory representations will be ‘rounded’ to the category exemplars provided. For example, even if a participant’s memory representation of the red target is no longer of the correct hue of red, their response will still be correct unless the representation deviates by a sufficient amount that it leads the participant to choose an entirely different category. This issue is especially prominent when using feature dimensions that might not fall into such distinct categories. For example, if participants are asked to remember subtly varying hues of a single colour (Olivers, Meijer, & Theeuwes, 2006), or orientations of lines (Vandenbroucke et al., 2011), their performance would depend partially on the response options presented during the test. For example, if during the response phase participants are asked to select the correct response from an array of choices, performance might differ significantly depending on whether the options are similar to each other (such as two subtly different shades of red) or dissimilar (such as when choosing between red and green), even if the quality of the actual memory representation itself was the same across both scenarios.

To capture this additional insight that is often lost using traditional paradigms, some researchers began to adapt paradigms also to measure the ‘precision’ of working memory.

These paradigms typically use a feature dimension that can be expressed on circular scale, mainly orientation (Bays & Husain, 2008), so that participants can be given a free choice of responses, rather than choosing between predetermined responses, as had been the case in previous experiments. Participants will typically be asked to remember one or more orientations (such as oriented bars or Gabor gratings) and then, when giving their response, will be asked to rotate a probe stimulus freely until it matches the orientation they recall. Their memory is then able to be measured as a function of the difference between the correct orientation and their response, with smaller deviations generally thought to reflect a better memory representation. This approach allows for researchers to measure the effects of their manipulations of interest on the precision of memory (Bays et al., 2009; Bays & Husain, 2008; Ma et al., 2014), often eliciting effects that would not be detected using traditional categorial designs.

Attention and Working Memory

Although the above review has included summaries of both attention and working memory as distinct mechanisms, they are strongly linked, to the extent that some researchers argue that working memory may be a form of attention, biasing the recurrent processing of internal representations over incoming external information (Chun, 2011). This may be thought of as a level of representation receiving inputs both from new processing of external information but also from feedback loops and recurrent processing of existing internal representations. If so, attention may operate between these two categories of inputs in a similar way as it does between multiple sources of external information and may act to bias one type of input relative to the other.

Hierarchical Representations

General

A full review of the structure of the visual perceptual system would be beyond the scope of this thesis. The key relevant point is that it is structured in a hierarchical manner (Schwartz & Giraldo, 2017), with levels of representation of particular forms of information progressing in a series of sequential and parallel stages. It is common for research to focus on a specific stage of processing at a time, but it important to remember that any such stage does not exist in isolation and likely has complex interactions with other stages, both up-stream and down-stream, via recurrent feed-forward and feed-back projections.

Hierarchical Representations in Perception/Consciousness

The existence of hierarchical representations is well acknowledged in research into conscious awareness, which often uses virtually identical paradigms to those in working memory (and may therefore often be measuring similar processes). For example, a distinction is typically made in studies of conscious awareness according to whether a participant is conscious of (able to report correctly) an object's identity or just its mere presence (Wyart & Tallon-Baudry, 2008). Even in the classic iconic memory experiments (Sperling, 1960), participants were able to recall the identities of only a few characters (those successfully encoded to working memory) but it was noted that they reported remembering other white letter-like stimuli, just not what the specific letters were. Although it is true that these participants were able to remember only a few letter identities correctly, it would be wrong to conclude that those letter identities were all that could be remembered at all. So entrenched was the model of working memory as a limited capacity of only a few high resolution items and nothing else, that some interpret such reports of vague additional information as evidence of having conscious experiences of information without being able to access or report it (phenomenological awareness; Block, 2011).

Such explanations overlook the fact that this additional information was indeed accessed and reported, but was not the type of information the study was designed to measure. It is probable that it would be impossible to find evidence that a participant was conscious of something that they could not report. Although participants may report that they were conscious of the letter identities but could not remember them, participants often make similar reports when the array includes non-letter symbols that share similar low-level properties as alphanumeric characters (de Gardelle, Sackur, & Kouider, 2009). This suggests that the participants' impressions that they were aware of all of the identities within the array may have been an illusion and they were actually aware of only other features (such as colour and location) or a 'coarse' representation of a 'letter-like' stimulus. This thesis focuses on this type of information, which is seemingly at a different 'level' of representation than the traditional high-fidelity individuated letter identity information that working memory paradigms are usually designed to measure.

Evidence suggests that these different levels may not necessarily be independent and parallel and that the formation of traditional high-resolution representations may proceed in a hierarchical manner, progressing from coarse to fine (Gao, Ding, Yang, Liang, & Shui, 2013). This process seems to progress from frontal brain areas towards posterior brain regions (Hochstein & Ahissar, 2002; Juan & Walsh, 2003). While this may initially appear counterintuitive, given this is the reverse direction that visual information is processed in (from 'early' posterior regions towards more anterior regions), it may not be so surprising when taking into consideration evidence of how robust representations (whether perceptual or memory) are formed. A full summary of this is beyond the scope and relevance of this description, but the key points are that feed-forward processing of visual information does not seem to be associated with conscious awareness and it is instead recurrent processing loops via feedback connections that lead to the formation of conscious representations (Boehler, Schoenfeld, Heinze, & Hopf, 2008; Lamme, 2001; Lamme & Roelfsema, 2000; Rees, 2007). Perhaps the key feature of such loops that seems to determine whether some information is conscious or not is whether such recurrent loops are 'stable' (Schurger, Sarigiannidis, Naccache, Sitt, & Dehaene, 2015). Given the importance in recurrent feedback projections in the formation of stable conscious percepts, the reverse hierarchy of perception is not so surprising.

The Association Between Consciousness and Working Memory

A key, often overlooked, feature of the above research into conscious perception is that the vast majority of such designs ask participants about their experiences only after the relevant stimuli have been removed and are no longer physically present. Therefore, regardless of what an individual may or may not have been conscious of while the stimuli were presented, all the evidence available is only that which can be maintained in working memory. The implication of this is that the considerations discussed above, such as reverse hierarchies and the need for stable recurrent processing, may be critical to working memory rather than consciousness itself. As a brief aside, although working memory and conscious awareness may be theoretically distinct, they may actually be the same fundamental mechanism. As already mentioned, the representations underlying working 'memory' form even while the stimuli are still physically present (Moore, 2001; Tsubomi et al., 2013; Vogel et al., 2006) and it is therefore likely that the core mechanism is the formation and maintenance of stable recurrent processing loops, modulated by attention such that they can persist over long periods and not be overwritten by intervening stimuli. The semantic distinction of whether such a representation is labelled as a conscious perception or working memory often pertains to whether the stimulus is still physically present or not. Note though that even this semantic distinction has potential flaws. A more parsimonious distinction may be that such a representation is classified as 'perception' if external information is being actively engaged, attended, and used to update or alter existing internal information, and classed as 'memory' when external information is being inhibited and the internal representation is being protected (however successfully), regardless of the physical presence or absence of the related stimulus.

Hierarchical Representations in Working Memory

The working memory literature has traditionally focused on the most advanced levels of representation, such as high precision representations of individuated objects. This may be partly because these types of stimuli are easier to quantify and report. For example, an experimenter can easily confirm whether a participant's report of an alphanumeric character, colour, orientation, shape, or photo identity matches the true original or not. However, as has been outlined above, neural representation involves multiple different levels of information, often in a hierarchical fashion. If working memory is largely mediated by attention biasing recurrent processing of internal representations over external incoming information, then it should be possible for internal representations to be maintained at any level with appropriate feedback connections to mediate such biasing.

The suggestion that working memory is not all or nothing (i.e. that one either remembers everything about an object or nothing) may seem uncontroversial as it is well understood that one might correctly remember one feature of an object (such as its colour) but misremember another feature (such as its shape; Fougnie & Alvarez, 2011). However, such research still often focusses on different feature dimensions at a similar 'level' or representation (traditional robust individuated working memory representations). This is still a valid point and an interesting avenue of research, but still limits its focus to only a narrow type of information and representations. Instead, evidence suggests that even if information measured in a traditional working memory task is not maintained robustly, other types of information at other 'levels' of representation may still be maintained (Nie, Muller, & Conci, 2017). The exact structure of these different levels of representation is nebulous and far beyond the scope of the current thesis, but a lot of this information appears hierarchical in nature, building from very coarse aggregate information calculated across an entire scene up to the traditional working memory representations of specific robust feature identities of individuated objects (Nie et al., 2017).

Regardless of the exact relationship between visual awareness and visual working memory, the paradigms used to study each are often indistinguishable and results from any such a study can be applied to both phenomena, regardless of which mechanism the authors were intending to measure. A key implication of this is that the above described findings of the reverse hierarchical nature of the formation of representations, and the importance of feedback projections for the formation and maintenance of stable recurrent processing, can also be applied to the study of working memory (Shin, Fabiani, & Gratton, 2006).

For example, when participants are able to report only 'coarse' information about a stimulus (such as that it was letter-like, or its approximate location), rather than reflecting different levels of awareness, such reports are more likely to reflect the maintenance of different levels of representation within working memory (Kouider, de Gardelle, Sackur, & Dupoux, 2010). Some have begun to suggest that the traditional narrow view of working memory mechanisms should be expanded to incorporate evidence of the maintenance of memory representations at hierarchical levels other than only the traditional most advanced stages (Brady & Alvarez, 2011; Brady, Konkle, & Alvarez, 2011; Ma et al., 2014; Suchow, Fougnie, Brady, & Alvarez, 2014).

It should be acknowledged that even within the main working memory field itself, some steps had already been taken towards changing concepts, such as the approach described above for measuring the precision of working memory rather than a correct/incorrect dichotomy (Bays et al., 2009). In principle such measures acknowledge that even if a memory representation may not be sufficient for a correct response at one level, this does not necessarily indicate that it was forgotten entirely and instead some other information may still have been maintained. However, such differences in precision do not necessarily map onto different hierarchical levels of representation, as even in conditions of low precision, the memory representations may still be advanced high-fidelity representations but of a very incorrect orientation. One approach to investigate whether different hierarchical levels may play a role in memory precision has been to ask participants explicitly to indicate not only the single specific orientation they recall, but also the range of other surrounding orientations they believe it might have been otherwise (Fougnie, Kanabar, Brady, & Alvarez, 2016). These confidence intervals were found not to be simply symmetrically centred on the main response, and instead showed an asymmetric bias towards what was the true correct orientation. This may indicate that, even if a participant has a high resolution memory of a specific but incorrect angle, other information at a coarser level of representation may be retained in the form of a probability distribution of values, revealing otherwise hidden information remembered regarding the correct orientation.

One of the key outstanding questions for hierarchical working memory representations is which stages of hierarchical processing can be maintained in memory. As described above, working memory relies on attentional mechanisms biasing internal recurrent processing, but it is unclear which levels of representation have the necessary feedback connections. It may be that some levels of representation, especially early levels, always reflect incoming information, if feedback connections are not present capable of sufficient biasing. An example of this might be that, given working memory and conscious perception appear tightly interlinked, it is difficult (arguably impossible) to block all conscious perception of incoming information in favour of focussing memory representations, which seems likely an evolutionary necessary limitation (as managing to block all incoming information would preclude the ability to respond to unexpected and potentially deadly events). Similarly, attention is not a binary process and even if some biasing mechanisms exist at a given level, it is unclear to what extent or with what resolution those attentional mechanisms may operate. It is possible that, though the feed-forward perceptual mechanisms are well understood for different levels of representation throughout the visual system, the feedback memory mechanisms that exist and the ways in which they operate and interact at these different levels, may be fundamentally different. This thesis will not aim to identify and map specific levels of hierarchical processing in memory but instead will use paradigms described below to investigate whether different conditions lead to observable behavioural changes which may involve broad underlying networks.

Group Level Information

So far, most of the research described has focussed on the memory of features of individual items. However, there are many other possible forms of information that might be maintained within working memory. One such example is relative encoding, where the critical information is the relative difference between two (or more) items rather than the absolute values of either (Clevenger & Hummel, 2014). For example, when trying to encode object locations, performance is greater when other nearby stimuli can be used as landmarks, allowing one to utilise knowledge of the relative positioning, than when needing to recall a target's absolute location in the absence of any such landmark cues (Aagten-Murphy & Bays, 2019). If such landmarks move, participants are also more accurate at reporting the target's location relative to the new landmark location than the absolute previous value independent of the landmark's move. Although this may seem obvious and intuitive, it is not easily accounted for by a model of working memory focussed only on independent representations of absolute values, which would predict the accuracy of absolute location memory should be unaffected by the presence of other items, or that the encoding of any such additional information should lead to interference and detrimental performance. It is worth noting that at a theoretical level, such relative encoding would not necessarily need to involve such specific information (such as the exact distance between two items) and could feasibly include coarser information, such as whether a stimulus is simply to the left of another (regardless of the exact position) or brighter (regardless of the exact difference in luminance).

Another example of information encoded across a group level is that of ensemble statistics, or summary statistics. It is not necessary, or even possible, to encode separately every minor detail of a surface or scene, such as the precise details of every leaf on a tree or blade of grass in a field. Instead, often complex information can be summarised approximately by their statistical features, such as in a manner analogous to how a Fourier Transformation may summarise a waveform in terms of its frequency spectrum. Examples of this might include texture perception (Dakin, 2014; Motoyoshi, Nishida, Sharan, & Adelson, 2007; Rosenholtz, 2014) or the segmentation of natural scenes in distinct objects or regions (Brady et al., 2017). It may be impossible to map every potential such statistical summary the brain employs, and many of the observable behavioural effects identified so far are likely mediated by a broad network of mechanisms. Some obvious such statistical properties already identified include the group's average value (Brady & Alvarez, 2011) and the variability (how homogenous the stimuli are; Norman, Heywood, & Kentridge, 2015).

However, the investigation of such mechanisms has traditionally focussed on their role in perception and their potential role in memory has been relatively neglected until recently. As previously noted, working memory can only operate at levels where feedback connections are sufficient to maintain stable internal representations, inhibiting new incoming information. Such mechanisms may operate effectively at some levels of processing and for some forms of such information, but may have limited or no control at others. As with other hierarchical information, one major difficulty is in how best to measure and quantify such effects. Perhaps the most studied is the ability to recall the average value of a group, even often in the absence of recalling reliably the identities of each specific item. Participants seem able to extract the mean size (Ariely, 2001; Chong & Treisman, 2003, 2005; Corbett & Melcher, 2014; Gorea et al., 2014), position (Alvarez & Oliva, 2008), colour (De Gardelle & Summerfield, 2011) and orientation (Solomon et al., 2016; Solomon, 2010).

Performance seems relatively unaffected by increases in the number of items in an array (Chong & Treisman, 2005; Haberman, Harp, & Whitney, 2009), suggesting that averaging mechanisms act in parallel across the entire array. Furthermore, separate averages appear to be able to be calculated both across the memory as a whole and distinct sub-groups (such as those sharing the same colour; Brady & Alvarez, 2011). However, costs from increased array sizes can be observed in some circumstances, such as when the individual features are more heterogenous (Marchant, Simons, & de Fockert, 2013). Evidence of the latter suggests that some form of serial mechanism may play a role, and that the apparently efficient averaging across even large groups may be facilitated by sub-sampling rather than necessarily processing every item (Gorea et al., 2014; Marchant et al., 2013). Such strategies cannot account for summary statistics mechanisms in general though. For example, representation of the variability of an array has shown to be sensitive to adaptation effects (a decrease in sensitivity from prolonged or repeated viewing), independent of individual or average values, suggesting a direct and automatic neural encoding (Norman et al., 2015). Furthermore, there is some evidence of performance sometimes increasing with greater numbers of stimuli (Robitaille & Harris, 2011), argued to be due to additional information for when calculating summary statistics, which would not be predicted by a strictly limited sub-sampling approach.

Regardless of the sampling approach, there is also debate regarding whether each item is weighted equally in any such calculations. With sufficiently large numbers of stimuli, it appears that items are weighted approximately uniformly (Juni, Singh, & Maloney, 2010), though it may be difficult with such large set sizes to detect potentially more subtle differences in weightings. With fewer items though, outliers seem be over-weighted when estimating the average location of dots (Moreland & Boynton, 2017) while outliers seem to be down-weighted when estimating average size or colour (De Gardelle & Summerfield, 2011). The difference here may be the relative difficulty in distinguishing distinct stimuli. When locating the average location of a cloud of dots, dots in dense clusters may be more difficult to distinguish, while when presented with a well-spaced array of circles, there is relatively less difficulty in perceiving each individual value. This might therefore indicate that outliers are down-weighted where possible, but clearer or more distinct stimuli may sometimes have a greater influence on such calculations.

Summary statistics are likely not always represented explicitly in perception, such that an individual can consciously access and report on the underlying processes reliably. However, evidence for sub-sampling strategies often comes from paradigms that explicitly require a response involving the summary statistic of interest, such as changing the size of a probe to indicate the estimated average size (Marchant et al., 2013). This type of paradigm might therefore encourage atypical purposeful sampling strategies, and so it may be preferable to use paradigms where the encoding and memory of ensemble statistics are observed indirectly. For example, even if every item in an array changes, detection of such a change is lower if summary statistics, such as the average value of each array, remains constant (Saiki & Holcombe, 2012; Ward, Bear, & Scholl, 2016). This finding is not easily accounted for by an effortful sub-sampling strategy. Similarly, participants appear able to detect changes in the summary statistics of an array even when they are unable to localise the changes of any specific individual items (Haberman & Whitney, 2011; Hollingworth, 2006). Furthermore, responses of the values of individual items in an array (as in traditional working memory paradigms) are often biased towards the group average (Brady & Alvarez, 2011; Corbett, 2017), suggesting that representations of the average cannot be accounted for simply as a purposefully calculated and encoded extra 'value' in working memory as other individual items. Taken together, the evidence is clear that though some situations may exist where group level information is artificially encoded and treated as though it were an additional individual value through purposeful strategic calculation and encoding, in many situations such ensemble encoding may play a similar role on memory as in perception, allowing for efficient representation of otherwise complex scenes and for the utility of information beyond that of individual objects.

Configuration Effect

The paradigm that is adopted in both Chapters 2 and 3 is the configuration effect paradigm (Jiang, Olson, & Chun, 2000). The configuration effect is most commonly demonstrated using the partial report paradigm in which one has to remember the identities of an array of stimuli but is subsequently tested on only one (or more) at random, not on every item. A baseline measure of traditional working memory capacity might be to present a probe at the same location as one of the original stimuli and ask the participant whether the feature identity of the probe is the same or different as the feature identity of that stimulus in the original memory array (change-detection paradigm; Gilchrist & Cowan, 2014; Gustavo Rohenkohl, Coull, & Nobre, 2011; Ilga G. Sligte et al., 2008; Geoffrey F Woodman, Vogel, & Luck, 2012). However, it has been observed that, if during the test phase a full array of stimuli is presented (but a response is still only required in relation to one cued probe), memory performance may be enhanced when the non-cued items in the probe array have the same identities as in the original memory array, or impaired if the non-cued items have different or altered characteristics (Boduroglu & Shah, 2006, 2009; Delvenne & Bruyer, 2006; Gmeindl, Nelson, Wiggin, & Reuter-Lorenz, 2011; Jiang et al., 2000; Mutluturk & Boduroglu, 2014; Papenmeier, Huff, & Schwan, 2012; Sanocki, Sellers, Mittelstadt, & Sulman, 2010; Silvis & Shapiro, 2014; Vidal, Gauchou, Tallon-Baudry, & O'Regan, 2005; Zimmermann, Schnier, & Lappe, 2010). In other words, relative to a single probe, memory performance can be aided by the presentation of the rest of the original memory array in the same configuration as it had originally been presented, or impaired by an explicit change in configuration.

If working memory were purely comprised of individual item representations then it might be expected that any decisions based on a single representation should be relatively unaffected by the presence or absence of additional independent information. It might even be expected that any such information at a critical period would lead to potential interference, distraction or confusion, rather than an improvement in performance. Importantly, therefore, these configuration effects can be taken to demonstrate that working memory representations of individual items are not fully independent of other items' representations. However, this effect is not universal and appears sensitive to various parameters. For example, a change in the spatial locations can disrupt memory performance for colour but the reverse does not appear to be true (Jiang et al., 2000). However, memory for location may not be totally independent of changes in other feature dimensions of the stimuli, as recent evidence has suggested location memory can be disrupted by changes in shape or orientation of placeholders (Toh, Sisk, & Jiang, 2020).

There are various possible explanations for the mechanisms underlying this pattern. One likely candidate may be the encoding of the 'average' value (Brady & Alvarez, 2011; Corbett & Melcher, 2014; Gorea et al., 2014; Solomon et al., 2016), as outlined earlier. If a participant manages to encode successfully only a sub-set of the original memory array into working memory, but the item that is probed is not one of those then, in the absence of any information from ensemble statistics, the participant would need merely to guess. However, if the participant has also encoded the 'average' value of the original memory array, this information may be utilised on trials where all the non-probed items are also re-presented in their original configuration. On such trials, if the 'average' value of the test array has differed from the remembered 'average' value of the original memory array, and it is known that all the non-probed items have their original value, then a participant could infer that the probed item must have therefore changed. On trials where the non-cued items are not presented, participants cannot leverage any changes in average values on trials where they have not remembered the probed item specifically, and on trials where other items are presented with different values, even though participants may know to ignore any changes or not in things such as average value, this additional interference may actively impair performance. This cannot be the sole mechanism behind the configuration effects though, as it would predict that performance would be unaffected by changes in the positions of items in the probe array so long as the same values are used in both arrays, but examples show that even in such situations, changes in configuration of only position can lead to similar improvements or detriments (Jiang et al., 2000).

Other candidate mechanisms for explaining configuration effects may include relative encoding, as outlined above, such that the relation between stimuli is encoded, potentially independently of the absolute values of either. This would still confer a similar benefit on some trials as remembering the average value would, such that even if the memory representation for the specific feature value of the probed item is poor, a change in the relation between the probed item and other items could still be sufficient to infer the correct response, but only if other items are re-presented in the same configuration as they had in the original memory array. Such information could not be utilised if the probe is presented in isolation or if a different configuration array is presented, in which case changes in relations between objects might confuse or interfere, leading to worse performance. Importantly, this mechanism would also lead to the observed effects even if the memory and probe arrays contain the same feature values just in different positions (Jiang et al., 2000).

It should be noted that such relative encoding need not act only between adjacent items and many complex interactions may be theoretically encoded across a scene. Furthermore, there may be many other similar mechanisms operating in various combinations. As noted earlier, one of the greatest challenges this avenue of research may face is the difficulty in defining, quantifying and reliably measuring the information of interest. Although work should continue to attempt to discern distinct mechanisms, it is not always necessary to understand the specific processes in a given instance for such effects to be informative. Investigating how observable phenomena such as the configuration effect are affected by carefully designed experimental manipulations may still provide important insights into the nature of the phenomena, even without being sure of the specific underlying neural mechanisms at this stage. This is the approach adopted throughout this thesis.

Far from epiphenomena existing only in specific artificial lab conditions, mechanisms such as hierarchical encoding and group-level ensemble encoding appear to be how the visual system naturally processes real world scenes and objects. Although such scenes and objects are more complex than simple colours or geometric shapes, the additional information available actually appears to enhance memory performance overall. Working memory capacity appears to be greater for real world objects than simple colours (Brady, Störmer, & Alvarez, 2016), and different levels or features of such structured objects appear to be forgotten independently rather than in an all-or-nothing manner (Brady, Konkle, & Alvarez, 2013). Similarly, real world scenes provide stronger context and configuration effects to aid change detection than similar scenes when inverted or scrambled (Zimmermann et al., 2010). Finally, studies of individual differences show that a participant's ability to utilise additional information can be independent for different levels of representation, but can also be stable within individuals (Haberman, Brady, & Alvarez, 2015). This suggests that the traditional advanced level of working memory information cannot be used as a simple proxy for an individual's ability to utilise information also at other levels of representation. It is clear that attempts to understand better how working memory may work in the real world would benefit from expanding their consideration of the capacity and form of working memory information from individual items toward structured representations (Brady et al., 2011; Brady & Tenenbaum, 2013; Cowan, Saults, & Clark, 2015; Jiang et al., 2000; Ma et al., 2014; Suchow et al., 2014).

The current thesis will focus on extending the understanding of hierarchical representations to consider other aspects of working memory for real world situations. Chapter 2 will still use basic abstract stimuli but where most such studies have used arrays of static, simultaneously presented stimuli, it will investigate how any such mechanisms may operate over time, as might occur in a complex, dynamic and changing environment of daily life. It will use sequentially presented stimuli to ask whether changes in temporal position have similar effects as changes in spatial position. Chapter 3 then seeks to extend existing findings from the visual domain to the tactile domain, motivated partially by the observation that texture perception requires integrating information over both space and time, and so potentially may show even greater hierarchical effects than visual perception.

Memory guided attention

The second broad theme within this thesis is that of memory guided attention. Often in life it is important that we can maintain relevant information 'in mind' for a short period even after it is no longer visible. Fortunately, we have evolved short-term memory mechanisms to enable this. Sometimes we can focus just on remembering the information (perhaps repeating it over in our heads) until we no longer require it, before then focussing on other things. However, often it is necessary to keep information in memory for later while focussing a secondary task in the meantime.

For example, after reading a sign approaching a roundabout for which exit to take, there are likely to be other similar sources of information (other signs with colours, words, orientations, etc.) visible in the field, perhaps even a repetition of the sign a little further along the road. It is here that issues may arise, if competing sources of information are presented in too similar a manner such that processing one may increase distraction by the other. Similarly, tasks often require us to remember information not because they physically vanish but because we need to maintain them across eye, head or body movements. Often the remembered information is therefore still physically present in the immediate environment and may even remain visible in peripheral vision while we fixate on the secondary task location. If this continued presence impairs allocation of attention towards the secondary task, such layouts might benefit from a voluntary occlusion of the original information (such that it can be refreshed if necessary but need not be present during performance of the secondary task, improving both).

An ongoing debate concerns how, in situations like this, the information held in visual working memory might affect, or be affected by, the secondary task, and how these effects differ when the information held in memory is relevant or irrelevant to the secondary task. This series of research aimed to build on existing data to investigate both how these interactions may occur for other fundamental feature dimensions and, importantly, how sensitive these interactions are to the similarity between the external and internal information.

Spatial Attention

A paradigm often used to research the interplay between attention and memory is the visual search task (Wolfe, 1994), where one has to search amongst an array of objects for one or more targets. A real-world example might include searching for a specific brand of cereal in a supermarket aisle. To achieve this, one must keep in memory what is being searched for (the search 'template') while attention moves around the scene and compares objects to the template until it finds a match. As outlined in the discussion, this is achieved by biasing the competition between incoming information such that some types of input are favoured or inhibited relative to others (biased competition model of attention; Desimone & Duncan, 1995). Depending on the conditions, this biasing may be able to guide attention in an efficient manner directly towards objects with features that match the template's (e.g. parallel search) or may need to perform a more effortful serial search (Treisman & Gelade, 1980).

However, in studies where the target of the search is the same as the item held in memory, it is unclear whether the memory representation itself is causing the attentional biasing or whether attentional search mechanisms are engaged independently of the template. In the latter case the information held in memory and the information that receives attentional biasing may correlate but may not reflect a truly causal relationship. To distinguish the effects of the search task demands from those of the working memory representation itself, one approach has been to adopt a dual task set-up where the memory representation is not relevant to the search task, to investigate whether merely holding information in memory may ‘automatically’ bias attention. Participants are asked to perform a typical memory task where they are asked to memorise some information then, after a delay period, perform a memory test. During the delay period, when in normal memory tests there may just be a blank delay or masking stimuli, participants are asked to complete a secondary task, usually a visual search task. The two tasks are kept independent from each other and the information relevant for the memory task is of no relevance for the search task.

A common example might be for the memory task to involve maintaining colour information while the search task would involve searching for a particular shape (Olivers et al., 2006). Critically, the stimuli in the search task might also include some colour information, usually in the form of a singleton (one item having a prominent colour while all other stimuli are monochromatic). The presence of such a salient distractor stimulus would be expected to capture some attention and therefore have a detrimental effect on the speed or efficiency of the shape search task. However, studies have repeatedly found that the relative magnitude of this distraction is greater when the singleton’s colour is the same as that in memory (Downing & Dodds, 2004; Olivers, 2009; Olivers & Eimer, 2011; Olivers et al., 2006; Soto, Heinke, Humphreys, & Blanco, 2005; Soto & Humphreys, 2007; Woodman & Luck, 2007). This is argued to indicate that holding the memory representation positively biases visual attention in favour of incoming information with features that match those in memory. Similar effects have also been observed where, rather than colour, the critical features were faces (Downing, 2000), line drawings of objects (Downing, 2000; Soto et al., 2005), shape (Downing & Dodds, 2004; Olivers et al., 2006; Soto & Humphreys, 2007, though Soto, Humphreys, & Heinke, 2006 failed to find such effects) and images of real world objects (Moores & Maxwell, 2008). However, the use of colour has generally been found to show more consistent and reliable effects than the other listed features (Soto et al., 2005) perhaps as it is a low level feature more effective at guiding attention, whereas many of the other examples require complex conjunctions of basic features. However, despite the above research, other studies have failed to find any such differences between when items in the search task share features with the memory items or not (Olivers, 2009), revealing that this effect does not generalise to all circumstances.

One consideration worth highlighting is that participants will likely adopt strategies to try to maximise their overall performance balanced across both tasks. They may therefore adopt a strategy that incurs a small cost to one of the tasks if it yields a greater improvement to the other, and a key consideration has been whether the above described between-task interactions could be accounted for in terms of adopting such strategies. Some studies use an exact match between the critical item in the memory and search task such that, even if its presence had no relevance to the search task, participants might have adopted a strategy to attend to the memory-matching face to 'refresh' their memory representation (e.g. Downing, 2000; Olivers, 2009) and improve performance on the memory task, given that doing so would have only a minor cost to the search identification task (on half of the trials it would mean attention was at the correct location while on the other half attention would be at the incorrect location, potentially leading to a net neutral effect). It is therefore difficult to determine whether these effects were due to the memories automatically biasing attention, or artificial strategies specific to those task parameters. This was further exacerbated by some studies inadvertently adopting designs that made the feature singleton informative for the search task such as by making it more likely to appear at the search target's location than at any given distractor location (Soto et al., 2005, 2006), making attending to and processing of the singleton advantageous to both tasks. Such task relevance would not itself account for the relative differences in performance depending on when the feature matches memory features or not, but it does preclude the interpretation that it is the memory representation itself that is biasing attention rather than purposeful strategic mechanisms. Other studies attempted to adopt the opposite approach and ensured that the feature singleton would never appear at the same location as the search target (Soto et al., 2005; Geoffrey Woodman & Luck, 2007), but this again leaves the singleton as relevant to the search task, as its presence could be used to avoid searching that location and therefore to guide attention to other locations. Indeed, such experiments found that when the feature was always at a non-target location, its presence (regardless of similarity to memory) actually improved search performance relative to trials where no singleton was present, and it was only when parameters were altered such that it could appear at any location that again its presence showed negative or distracting effects on the search task (Woodman & Luck, 2007). When the memory-matching item was truly uninformative for the search task, search task reaction times were still slowed by its presence relative to an item that did not match memory contents (e.g. Zhang, Zhang, Huang, Kong, & Wang, 2011), confirming that such strategies cannot account fully for such memory guidance, but highlight that research should be careful to avoid such confounds.

It could still be argued that, though the above approach does introduce a cost to attending to the memory-matching item, participants may feel this cost is worth it relative to the potential gains from refreshing the memory item, such that they would still make a net improvement to performance. This is especially true for studies where the critical stimulus is an exact match between the memory task and search task (e.g. in Olivers et al., 2006). To account for this, some have used ‘memory-matching’ items that are similar, but never identical, to the memory item (such as a different shade of red). The same pattern of results is still observed, even when potential gains from memory refreshing are minimised in this way (e.g. Olivers, 2009; Olivers et al., 2006), suggesting that observed memory guided effects cannot be accounted for purely be strategic refreshing of memory-matching items, though again research should take care to avoid encouraging such confounding strategies in their designs.

Although the previously reported memory guided attention effects persist even after the removal of confounds in the design that might encourage artificial strategies for the deployment of attention, other key parameters have been found to be critical for the effect. Evidence of such memory guidance has been found most reliably when only one item is being held in memory (Houtkamp & Roelfsema, 2006; van Moorselaar, Theeuwes, & Olivers, 2014), and is often not observed when more than one item needs to be held in memory across a given trial (Soto & Humphreys, 2008; Woodman, Vogel, & Luck, 2001). There are some exceptions to this and some studies do report positive effects even when holding two items in memory, but not with larger memory loads (Moores & Maxwell, 2008; Zhang et al., 2011). Importantly, it does not appear to be the number of items in memory per se, but the ‘state’ of such internal representations. For example, if two items are initially encoded but one of those items is later designated the sole memory item via a retro cue (such that other item can be forgotten and is of no further relevance) then the cued item does still exhibit memory guided attention, though not if no cue is presented and both items are relevant (van Moorselaar et al., 2014). Even on trials where two memory items are maintained throughout, if the order in which the two items will be tested is known by the participant in advance, the item that will be tested first does lead to memory guided attention but not the second memory item, and no such guidance occurs if the order of memory item testing is unknown (van Moorselaar et al., 2014). These results reveal that the lack of memory guided attention effects cannot be accounted for by difficulties in encoding multiple items relative to encoding only a single item, and instead relate to the immediate relevance of a given item in memory.

Together, such evidence has been taken to indicate that there are different ‘states’ within working memory, a single active search template (which can bias attention) and dormant accessory items (which, though still in memory, do not bias attention; Olivers, Peters, Houtkamp, & Roelfsema, 2011; Peters, Goebel, & Roelfsema, 2009; van Moorselaar et al., 2014). Only when a single clear item is held in the search template (either when only a single item is in memory, or when a specific memory item is designated as the ‘active’ representation) is biasing of visual attention reliably observed. Surprisingly, this seems specific to when the memory items are irrelevant to the search task, as evidence suggests that visual search can be guided by two simultaneous search targets (each needing to be maintained in working memory; Bahle, Thayer, Mordkoff, & Hollingworth, 2019).

Related to this, memory guided attention effects are often not observed when the target in the search task can change on each trial (Downing & Dodds, 2004; Houtkamp & Roelfsema, 2006; Olivers, 2009), such as if the specific feature to search for is revealed to the participant only at the start of each trial. This again has been interpreted as being due to only a single active search template being able to guide attention. If the search task changes on every trial then each target must be held temporarily in working memory and would therefore occupy the active template slot, preventing attentional biasing from any additional ‘accessory’ working memory items. However, when the search task target is consistent across trials then it can eventually be offloaded to long-term memory, allowing for another item, the item in working memory, to occupy the active template slot instead and subsequently bias attention (Olivers et al., 2011; van Moorselaar et al., 2014).

It does appear to be the case though that such guidance is from working memory itself. If items are merely presented before the search task but no memorisation is required then this prime typically does not appear to guide attention in the same way as when memorised (Soto et al., 2005). Similarly, if a previously encoded memory item is tested prior to the search task then no subsequent attentional guidance effects are observed (Olivers & Eimer, 2011; Olivers et al., 2006). Furthermore, only the memory relevant information shows evidence of biasing attention. If the encoded memory items also had irrelevant secondary features (such as shape), no evidence of guidance from those features are observed (Olivers et al., 2006).

Many questions remain outstanding though. The focus of this thesis is how any such memory biasing of attention operates at different levels of representation. Although it is established that a feature in memory can bias attention towards other similar feature values that aren't necessarily identical (Olivers et al., 2006), little is understood about how the magnitude of such mechanisms vary as a function of the similarity between the memory item and search task singleton. Two key possibilities might be: 1) that the magnitude of such biasing effects gradually decreases with increasing levels of difference between the specific value in memory and the specific value of the singleton in the search task; or 2) that, if memory guided attention occurs at a coarser level of representation that represents a range of values (such as all shades of red), then this might lead to a similar biasing for all other values in the range represented by that channel or level of representation. All previous studies outlined above have used a binary categorisation of search target singletons as either the same (or similar) to the memory item or different. However, if a continuous measure of similarity were adopted then this may enable more subtle effects to be observed. For example, if two stimuli are held in memory, previous results have mostly found no reliable differences between when one of the values was present in the search task or a different value trial (Soto & Humphreys, 2008; Woodman et al., 2001). However, a continuous measure would allow for the similarity between the two exemplars chosen, and between the feature singleton in the memory array to be better controlled. For example, when holding in memory a red and yellow item, it remains possible that an orange stimulus might show greater memory biasing than a blue stimulus. Such effects would be unlikely to be detected by previous paradigms though as the relative differences between the multiple memory items were chosen randomly, and the extent to which the 'different' item actually differed from the others was poorly controlled.

A major problem with attempting to develop an approach with a more continuous measure of similarity is that the effect has been found most reliably using colour. As outlined earlier, this may be because many of the other features used (faces, shapes, complex images) constitute higher level representations that require often complex conjunctions of many features, while colour is a low-level basic feature, potentially more efficient in biasing attention across the visual field. However, colour is difficult to represent on a continuous measure. Therefore, the first aim of Chapter 4 is to investigate whether previous memory guided attention effects can also be replicated with other low level feature categories that may lend themselves better as continuous measures (specifically, angle of orientation and angle of direction of motion).

Temporal Attention

All of the research described so far has focussed on the biasing of spatial attention, such that, when multiple stimuli are presented simultaneously and are therefore competing for limited representation resources, the contents of memory may bias processing in favour of some items over others. However, it has long been recognised that attention operates over time as well as space (Griffin et al., 2001). For example, one's level of alertness may fluctuate moment to moment, at one point focussed and poised to perceive and response to incoming input, while at other times distracted or daydreaming, potentially missing presented stimuli or being sluggish to respond. For example, trial-by-trial variability of occipital alpha (a neural measure considered to be negatively correlated with attention to external stimulation) was found to correlate negatively with task performance on a perceptual task (Macdonald et al., 2011). Beyond random variability, temporal attention can also vary in a purposeful and systematic manner. For example, when monitoring for the onset of a target stimulus, the level of attention and vigilance may vary between time periods where it is highly likely or unlikely to appear (Coull & Nobre, 1998).

In addition to temporal variations of attention affecting how the onset of new stimulation is processed, there are also limitations to the temporal 'precision' with which we can represent information. This is perhaps most easily demonstrated perceptually by a stimulus that is flickering rapidly. If the gap between the stimulus presentations is long enough then one will perceive it correctly as a flickering stimulus, oscillating between present and absent. However, as the gap between presentations gets shorter, a threshold is reached at which the stimulus appears to remain constant (Kietzman & Sutton, 1968). The most familiar example of this is probably the old CRT monitors and televisions, which only ever had one pixel illuminated at any given moment but appeared to users as a stable and constant image. This may reflect the physical limitations of the perceptual system, but there also appear to be analogous temporal limitations to attention and neural representation, such that there is a limit to how much information that can be processed effectively to a given level of representation within a given period of time (Busch & VanRullen, 2010; Macdonald, Cavanagh, & VanRullen, 2014; Samaha & Postle, 2015; VanRullen, 2016).

Some argue that the brain might perceive the world in a series of fluctuating or discrete samples (Macdonald et al., 2014; VanRullen, 2016), arguably linked to an individual's peak alpha frequency (PAF) at rest (Samaha & Postle, 2015). One's PAF is defined as the frequency in the alpha band (8-12Hz) with the greatest power while at rest. This peak frequency has been found to have high stability over time (Vanrullen & MacDonald, 2012) and previous research has found that an individual's PAF positively correlates with measures of perceptual temporal precision (e.g. Samaha & Postle, 2015), or negatively correlated with the precision of representation of a dynamically moving stimulus (e.g. Howard, Arnold, & Belmonte, 2017).

It is thought that this discrete sampling might reflect the time taken to process and a 'snapshot' of information and integrate it into a stable representation through recurrent processing. However, as discussed earlier, different levels of representation may behave differently. For example, it may be that such sampling only occurs at later, more advanced stages of representation where information needs to be integrated from many processes and form explicit representation, whereas earlier levels of representation such as simple detection may have much higher levels of resolution (such as the physical limits as described earlier). Similarly, the evidence does not necessarily require that such sampling is truly discrete, and it may be that the relative biasing given to internal and external inputs may oscillate over time, but salient stimuli may still be able to 'break through' even if they are briefly presented between snapshots.

A more established example of the temporal limits of attention is the attentional blink (AB) paradigm (Raymond, Shapiro, & Arnell, 1992). This involves presenting participants with a rapid stream of stimuli (RSVP, typically 10Hz), one at a time at the same spatial location. Embedded within this stream are two 'targets' that need to be remembered and recalled at the end of the trial. Typically, participants have high accuracy for reporting the first target (T1), but accuracy for the second target (T2) depends on the temporal interval between the T1 and T2. When this interval is sufficiently short, there is a marked decrease in T2 accuracy, with this period termed the 'attentional blink'. Traditionally this is considered to occur because it would take time for T1 to be fully processed and encoded to a stable memory representation, and further processing of potentially distracting information is inhibited during this period (Chun & Potter, 1995), and it has been suggested this might also have a link to occipital alpha frequencies (Shapiro, Hanslmayr, Enns, & Lleras, 2017; Zauner et al., 2012). Therefore, if T2 is presented while resources are still taken up encoding T1, T2 may be 'missed'.

However, encoding of T1 into memory is not the only suggested cause of the AB. The Boost-Bounce hypothesis (Olivers & Meeter, 2008) derives from evidence that if only T1 and T2 are presented (with no non-target presented in the intervening period), little or no AB is observed. If the effect were solely due to the time taken to encode T1 into working memory before it can process further stimuli, an AB should be observed in the same time period regardless of what stimuli are presented in the meantime. Furthermore, if the blank interval between T1 and the first non-target following T1 (D1, or distractor 1) is extended slightly therefore delaying the onset of D1, the temporal profile of the subsequent blink also appears to follow a similar delay (Olivers & Meeter, 2008). This suggests that the blink may be time-locked to the onset of D1 rather than to the onset of T1 itself, and that it may not be only encoding of T1 that causes the AB, but the subsequent interference from the first subsequent distractor (D1). The Boost-Bounce theory explains this as T1 triggering an increase in attention. However, when this increase of attention then processes D1, which is recognised as a distractor, the visual system responds by triggering inhibition of incoming information. It is suggested that this inhibition may overcompensate and cause an 'overshoot', resulting in a period where incoming sensory information is inhibited relative to the baseline, impairing processing of T2, and causing the attentional blink.

Taking the evidence together, both mechanisms may play some role, but it is important to note that though most studies will state that T2 is not detected, the task typically involves identification rather than true detection. The evidence suggests that participants are not able to identify what letter or number T2 is, rather than not being able to perceive that there is a stimulus there at all (Nieuwenhuis & de Kleijn, 2011; though, Sergent & Dehaene, 2004 suggest even simple detection of T2 may indeed be inhibited). This again is important as different levels of representation may have different temporal properties, and letter identification may be affected while detection mechanisms are not.

It remains unclear though whether the contents of working memory may affect temporal as well as spatial patterns of attentional allocation. Some studies have suggested that inhibition of external information increases during the maintenance period of a working memory task, to protect the internal representation from erroneous integration of distracting information (Chun, 2011). However, this is primarily when the memory maintenance is the sole task. As with the spatial attention studies described above, it is less clear how maintaining information in memory may affect the temporal processing of subsequent information when performing a secondary task. For example, the mechanisms of memory maintenance might affect the temporal 'sampling' of the world, or similarities between the remembered information and incoming information might affect the rate at which it is processed. The focus of Chapter 5 is therefore to investigate whether the contents of memory might affect the temporal processing of information, primarily by adapting the dual task paradigm used in the spatial attention experiments described above. Instead of a spatial search task during the memory maintenance period, participants will perform an AB task, to investigate whether the magnitude of the ’blink’ varies depending on the presence or absence of featural information matching that held in memory.

Summary

As outlined at the start of this chapter, the core theme of this thesis is the interaction between visual attention and visual working memory. More specifically, this thesis aims to build on recent literature demonstrating that traditional views of working memory may be too narrow and that it may instead be structured as a multi-level hierarchy. The first core questions addressed here are how information across multiple items, beyond each item’s individual identity, is processed across time (Chapter 2), or in the tactile domain (Chapter 3). The thesis then aims to investigate how information held in working memory may affect the subsequent processing of subsequent information, with particular focus on how information at different levels of representation of a hierarchically structured working memory may interact (Chapters 4 and 5).

Chapter 2: Temporal Hierarchical Representations and Working Memory Encoding

General Introduction

As outlined in the previous chapter, there is a much wider range of information that can be maintained in working memory than the high-resolution and specific feature values of individuated objects that are the focus of the vast majority of traditional working memory research. For example, one might remember the average feature value of a group of stimuli even in the absence of memories of every individual item (Brady & Alvarez, 2011; J. A. Solomon et al., 2016), one might remember the relative distance between two targets even with a poor memory for the absolute positions of either (Aagten-Murphy & Bays, 2019), or even if one is not sure of the specific correct identity, one might have some idea of which range of values it was more likely to be than others (Fougnie et al., 2016). In light of this, there have been calls to expand existing concepts of working memory to account for the growing evidence of these other forms of information (Brady et al., 2011; Ma et al., 2014; Suchow et al., 2014).

It can be challenging to investigate mechanisms such as these though. Such processes or levels of representation can be difficult to articulate, quantify and measure reliably and although attempts should continue in mapping out specific effects, another approach has been to develop paradigms that can be used to detect changes in different conditions that could be me mediated by a broad combination of unknown mechanisms. This still allows for broad interpretations to be drawn on how hierarchical representations operate in working memory in general, without necessarily needing to isolate specific aspects. One such paradigm, and the paradigm that will be used in both this chapter and the next, is the configuration effect paradigm (Jiang et al., 2000).

As described in the previous chapter, the key feature of the configuration effect paradigm is that memory performance for an individual item is affected by the presence of other information during the memory test phase. If the items other than the probed target are presented identical as they were in the original encoding memory array then memory performance for the cued target is enhanced relative to when the probe is presented in isolation, and if new or different stimuli are presented then target performance becomes worse (Boduroglu & Shah, 2006, 2009; Delvenne & Bruyer, 2006; Gmeindl et al., 2011; Jiang et al., 2000; Mutluturk & Boduroglu, 2014; Papenmeier et al., 2012; Sanocki et al., 2010; Silvis & Shapiro, 2014; Vidal et al., 2005; Zimmermann et al., 2010).

There are various possible mechanisms that might contribute towards this effect. If a participant remembers that one of the items was yellow but not necessarily which one, then revealing the identities of the non-cued items may help to resolve this uncertainty (confirming whether the yellow one was a different item or not). However, this mis-binding is likely to occur only when requiring a combination of two features (such as colour and location), yet the configuration changes occur even when the task is unidimensional (spatial location; Jiang et al., 2000). Instead, other likely mechanisms include summary statistics, such as average value (Brady & Alvarez, 2011; Corbett & Melcher, 2014; Gorea et al., 2014; Solomon et al., 2016), or relative encoding between stimuli (Clevenger & Hummel, 2014). Regardless of the specific combination of mechanisms, the key is whether any group level information encoded during the initial memory array can be utilised during the memory test. For example, if a participant were to encode some information such as a group level mean, or the relative difference between the target and an adjacent stimulus, then even if they could not remember the specific value of the probed target at all, they could still detect a change in the array if the average value changed or the relation between the target and an adjacent stimulus had changed. However, for participants to be able to utilise changes group level statistics to infer changes of the individual target requires that the other items remain the same during both the encoding and test phases. If the non-targets change between presentations then it would be more difficult to distinguish whether the resultant changes in group level information also included a change in the target item or not. Similarly, if the probe is presented in isolation then such information cannot be utilised at all and memory performance would depend on only the traditional working memory representation for that stimulus.

Most research into group level mechanisms such as ensemble encoding or relative encoding has focussed on static displays where all of the relevant information is presented simultaneously. However, in daily life the information we are presented with is often dynamic and ever-changing. Therefore, a key outstanding question is how such representations operate over time rather than only over space.

One potential mechanism that might underly some of the observed effects such as calculation of group averages is that of receptive fields. These are broadly the regions of space across which a neuron is sensitive to, with a small receptive field meaning that neuron codes for information only in a very small specific region of space. Neurons at different levels within the visual system often have different size receptive fields, with especially some of the later stages able to integrate information from multiple objects across a wide region of space (Hochstein & Ahissar, 2002). Evidence of integrating and averaging inputs over time does exist, but these are typically at much shorter timescales and focus on the limits of whether stimuli are perceived as separate or as one continuous stimulus (Kietzman & Sutton, 1968). This may not be analogous to the effects observed in spatial paradigms, and it remains ambiguous whether similar effects might occur at longer timescales. Another reason for suspecting that such mechanisms may not necessarily generalise to sequential information is that the encoding of the relative difference between two simultaneously present stimuli could in theory operate independently of mechanisms for encoding the specific values of either individual object (for example, the space between two adjacent objects could be perceived and encoded as its own distinct value, without reference to the absolute locations of either object). However, it is not clear how this could be possible for sequentially presented stimuli as the specific values of each object would need to be maintained and encoded in memory for any relative information to be compared with subsequently presented stimuli.

These considerations might predict that therefore such effects simply would not be possible for sequentially presented stimuli, yet evidence suggests that at least some group-level information is encoded across temporally distinct stimuli. When two arrays are presented in quick succession, it appears that if the interval between them is very short (less than 500ms) then contextual information may be integrated across the two, whereas if the interval is longer then contextual information from each array may have distinct influences (Jiang & Kumar, 2004). However, even at slightly longer intervals (though still less than one second) Gestalt-like principles (such as grouping based on similarity) have been observed even for sequentially presented items (Gao, Gao, Tang, Shui, & Shen, 2015). The relative locations of objects also appears to be encoded for items presented sequentially one at a time (Ryan & Villate, 2009).

Evidence also suggests at least some ensemble encoding such as ‘average’ occurs even across temporally separate stimuli. For example, participants seem able to identify the ‘average’ size (Gorea et al., 2014) and emotional expression (Haberman et al., 2009) across stimuli presented sequentially, potentially even without necessarily encoding each individual item (Corbett & Oriet, 2011). It is worth noting that these temporal effects do not necessarily need to be strictly group-level representations and such averaging has also been shown to occur for a single stimulus smoothly expanding and contracting over time (Albrecht & Scholl, 2010). As with calculating summary statistics over space, it is possible that such calculations do not weight every item equally. For example, in the dynamically changing single stimulus, moments of expansion seemed to bias estimates of the average more than moments of contraction (Albrecht & Scholl, 2010). Generally it seems that with sufficiently large set sizes, a recency effect can be observed, such that items towards the end of a stream have a greater influence on estimates of average value then earlier items (Hubert-wallander & Boynton, 2015; Kool, Conway, & Turk-Browne, 2014). However, the reverse appears true for average location, and earlier items seem to have a greater influence than later items (Hubert-wallander & Boynton, 2015). This discrepancy could reflect also the findings that when items are presented simultaneously, that estimates of average location are biased more by those items in less dense regions than those in more dense regions (Moreland & Boynton, 2017). If the locations of subsequently presented items are integrated in some form, the earlier items could be considered to be in less dense regions as, at the time they were presented there were fewer surrounding competing items than for the later items.

However, there are potential issues with interpreting such 'averaging'. It is theoretically possible to calculate the average value of a sequential stream just by maintaining a running average and a count of the number of stimuli and updating the running average based on each new stimulus but with a progressively reduced weighting. This would allow participants to calculate an accurate average value by needing to maintain only these two values, without needing to encode any of the actual individual values. Indeed, evidence seems to suggest that the ability to calculate the average value in a sequential stream of items does not depend on accurate encoding of each individual item (Corbett & Oriet, 2011). Whether participants did, or even could, use such a strategy is unclear, but is a particular concern as the above studies used paradigms that required the average value to be considered and reported explicitly, therefore increasing the chance of artificial strategies. Such an issue would not apply to all identified forms of ensemble encoding and, for example, no such strategy seems likely for measures of the distribution of values (variability) within a sequential stream. However, it is also unclear how one might measure participants' representations of the distribution of values explicitly. The studies in this chapter will therefore adopt the configuration effect paradigm (Jiang et al., 2000), as this has the advantage of using an indirect measure to detect any ensemble encoding differences with a reduced risk of participants adopting artificial explicit strategies.

Experiment 1: Configuration Effects for Spatially and Sequentially Presented Stimuli

Introduction

The aim of this experiment is to investigate whether whatever mechanisms give rise to the configuration effect (Jiang et al., 2000) across space when using arrays of simultaneously presented stimuli also operate across time when using arrays of sequentially presented stimuli. There have been previous attempts to investigate temporal configuration effects. Perhaps the most relevant example presented participants with an array of shapes, either sequentially or simultaneously, and though overall memory performance was higher when the items were presented simultaneously than sequentially, both forms of presentation showed apparent sensitivity to changes in configuration (Blalock & Clegg, 2010). However, there are some issues with the design that make the results ambiguous to interpret. The task in this example was to detect any change in the array, and in the critical condition where the identities of the items changed, only two of them swapped, meaning that the other half of the array was identical. Furthermore, performance in this condition when the initial array had been presented sequentially was not distinguishable from chance. Therefore, rather than evidence of ensemble encoding across sequential stimuli, it is possible that such group level information was not well encoded at all and responses relied instead on detecting changes in individual items. It is not clear why this would have led to such poor performance, but the relative difference could be accounted for by a bias towards a ‘no change’ response (as at least half of the items would still be in the same location) which would lead to a greater proportion correct on the ‘no change’ than ‘change’ trials. A measure of sensitivity (d’) was included but only reported for comparisons between simultaneous and sequential conditions, not between the different configuration conditions within either presentation type. The stimuli were also presented for relatively long durations, of one second per stimulus (or four seconds if simultaneous), reducing the likelihood of robust natural temporal integration over such long time periods. The current study will require participants to detect the change of only a single cued target item, rather than from the entire array, such that the properties of the non-target stimuli should not affect any purposeful response strategies to the same degree, especially as on trials where the configuration changes, every item will change rather than only a sub-set. Sensitivity measures will also be included for all comparisons.

In another study (Olivers & Schreij, 2014), participants performed a typical working memory task where they would be presented with an array of colours to maintain in memory across a delay period where the items are not visible, before being shown a probe array and being asked to indicate whether any of the items in the second array were different than in the original array. Critically though, these arrays did not have abrupt onsets and offsets as is often the case with visual stimuli, and instead the arrays would rapidly move onto or off of the screen from the edges of the screen. Interestingly, performance for the memory task when the direction that the encoding array and the test array appeared from was consistent, relative to when the test array appeared on screen from a different direction to the original array. This is a surprising result as it is not easily explained using typical ensemble encoding mechanisms. The colour information in the array itself is unaffected by the direction it enters from, and each transition is rapid (in motion for only 150ms), limiting the likelihood of meaningful processing occurring during the motion itself. However, the result appears to indicate that objects are not encoded in a purely spatial manner but are instead encoded along with additional information from shortly before or after. The current experiment will aim to build on these results by presenting items in an array sequentially, and then keeping the information just before or after a critical moment either consistent or different between the encoding and test array. Importantly, this change of information would relate to the critical feature of the array itself, rather than irrelevant motion information as in the above design, hopefully demonstrating such temporal integration effects more robustly.

In a third design, participants are shown two groups, each of four stimuli, either presented one group at a time, or both groups simultaneously. When the test array was of only one of these sub-groups, memory performance was higher when the original encoding array had also been presented as two separate arrays (Ihssen, Linden, & Shapiro, 2010; Silvis & Shapiro, 2014), and when the original memory array was presented as all eight items simultaneously, performance was better if the test array also presented all eight items than if presenting only one sub-group of four items (Silvis & Shapiro, 2014). The initial explanation for these results could be the consistency of the item presentation across the two arrays allowing for improved use of any encoded summary statistics, and would suggest that such statistics are calculated mainly across simultaneous input and are not similarly combined with temporally distinct inputs. However, although there was an improvement when both memory and test arrays presented the two groups separately, there was a similar improvement if the encoding array presented all eight items simultaneously but twice (Ihssen et al., 2010), suggesting other factors such as consolidation time could instead account for the observed effects. Also, memory performance was similar when the test array presented all eight items simultaneously, regardless of what type of encoding array presentation type was used (Silvis

& Shapiro, 2014). However, there is a potential confound of task difficulty when having to detect a change amongst only four stimuli than amongst eight, which may have cancelled out the expected improvement when the test array was presented simultaneously. Also, the two arrays of items had different critical feature dimensions (one was of four colours and the other was of four shapes), meaning that there would be minimal effects from ensemble encoding across the two arrays even when simultaneous, and any such effects would be even less likely for any integration across temporally distinct arrays. This difference between the two arrays might also contribute towards why memory was high for a four item test array even when the encoding array contained all eight items but presented twice, as participants may have adopted a strategy to focus on one array in the first presentation and another array in the latter, perhaps reducing further what little between-array ensemble encoding may have occurred, and therefore making the condition closer to when the two arrays were presented separately. Finally, the test was to detect if any of the items in an array changed, rather than a specific cued item as in other forms of the configuration effect paradigm (Vidal et al., 2005). This means that even when one item did change, the other three items would always be the same in both the memory and test arrays. This explicit focus on group level information and the lack of meaningful change in configuration might again have encouraged purposeful encoding strategies.

The current experiment will therefore attempt to address some of these identified issues. Each array will consist of five coloured circles and just before the test array is presented, one item will be cued as the target, with the task to identify whether that colour changes between the encoding and test array, regardless of the other items. This should keep the core task a consistent difficulty comparable as, no matter the presentation type, the task should always be to remember the five colours in the encoding array, and to report the identity of only one of them (rather than needing sometimes to detect a change in four or eight items, as in the above design (Silvis & Shapiro, 2014). Furthermore, the focus of the task on remembering the specific identities of each item should hopefully minimise the incentive for participants to adopt abnormal encoding strategies and so any observed configuration effects should reflect more natural or realistic mechanisms. Each array will be presented either fully simultaneously or fully sequentially, and every item will have the same task relevant factor (colour). This should allow for the detection of any configural information being integrated over time. In the studies outlined above, the test array was always presented simultaneously, so the inclusion of trials where also the test array has a sequential presentation should allow us to measure better any sensitivities in changes to the temporal as well as spatial properties of the stimuli. Finally, the use of a cued target means that the non-cued items can be either kept the same as during encoding or have a different configuration, allowing for a more direct comparison of changes in spatial configuration while keeping other parameters consistent.

Importantly, the design includes trials where the arrays are presented simultaneously, as if all arrays were presented only sequentially then a null result would be difficult to interpret. It would be unclear whether such a lack of effect indicates that no configural effects occur over temporally distinct stimuli, or whether there was some other issue with the parameters or stimuli used that might interfere with even previously established spatial configuration effects. The inclusion of both types of presentation types allows for a more reliable comparison between simultaneous and sequential information using otherwise identical parameters.

There are many potential questions that could be asked of temporal effects on configuration effect mechanisms. For example, even if the spatial information of an array was kept constant, whether keeping the order that they appear consistent between the encoding and test arrays may make a difference. Similarly, if cueing a single item as the target, whether the magnitude of configuration effects differ when the changes are to items temporally close to the target versus other items in the array presented temporally earlier or later. However, it was not possible to test all of these questions immediately as the number of conditions would require the testing session to be unfeasibly long. Therefore, in the current experiment, spatial and temporal information will be varied together, such that items in the encoding and test arrays have either identical spatial and temporal parameters, or every item (other than potentially the target) has a different spatial and temporal value. It is hoped that by using instances of maximal similarity or maximal difference should increase sensitivity to detecting any temporal effects and, if found, could then be explored further by more complex or subtle design manipulations.

It is predicted that memory performance will be greater when the non-cued items have the same spatial configuration versus when they have a different configuration, with performance on trials with only an isolated probe stimulus (no other array items) somewhere between the two. It is also predicted that performance will be greater when the encoding and probe arrays are both presented either sequentially or simultaneously, and will be impaired if presentation method swaps between phases.

Method

Participants

Twenty participants volunteered in return for course credit. This was comparable to, and often exceeded, many previous studies showing robust configuration effects (e.g. 8-20 participants; Boduroglu & Shah, 2006; Papenmeier et al., 2012; Jiang et al., 2000; Vidal et al., 2005), including those showing potential temporal-like effects (e.g. 12-14 participants, Silvis & Shapiro, 2014; Olivers & Schreij, 2014). Demographic information was not recorded for three participants due to technical error, but the remaining seventeen were all female, with ages ranging from 18 to 36 (M=19.59, SD= 4.29), and all reported normal or corrected to normal vision. All procedures were reviewed and approved by the Departmental Ethics Committee.

Stimuli and Apparatus

All stimuli were presented on a CRT monitor driven at 85Hz, at an approximate viewing distance of 57cm. All responses were made using a keyboard by pressing either the left or right arrow keys to indicate 'same' or 'change' trials (which key corresponded to which response was counterbalanced). The experiment was created and run using PsychoPy (Peirce et al., 2019).

All memory stimuli were solid-coloured discs (1.1° diameter). These could be one of eight colours (every combination of each RGB phosphor set either to maximum or minimum). There could be no more than one disc of each colour in each array. A mid-grey background was displayed throughout the experiment.

The memory array was composed of five discs arranged equally spaced on the circumference of an invisible circle with a radius of 1.6°, with one item positioned at the lowest point, at the vertical meridian. On some trials each item of the array was presented sequentially in a random order, while on other trials all items were presented simultaneously.

The memory cue consisted of a 0.96° long black line, starting 0.16° from fixation, and pointing from the centre of the screen to the previous position of the target item.

The test array always presented a memory probe disc at the location of the cued target. On 50% trials, this disc could be either the same colour as the original target item (target-same trial), while on the other 50% it was a new colour that had not been present in the original memory array (target-change trial).

On 80% of trials, discs were also presented at all other locations (config trials), while on 20% of trials no other discs were present (isolation trials). When the other discs were presented, they could again be presented either simultaneously or sequentially. On each config trial, either every non-cued colour from the memory array was presented at the same corresponding locations in the probe array (config-same, 40% of trials), or all at different locations (config-change 40% of trials).

The presentation method (simultaneous versus sequential) of the test array on config trials could be either the same or different as the presentation method of the memory array, with equal probability. On trials where both arrays used sequential presentation, config-same trials kept the same order of presentation for both arrays, while config-change trials presented each colour and location in a different temporal position.

At all times during the experiment, a small white fixation point (0.2° diameter) was visible. This was a solid white disc from 705ms prior to the memory array onset, until the end of presentation of the probe array. Once the probe array finished, it would change to a hollow white ring.

Procedure

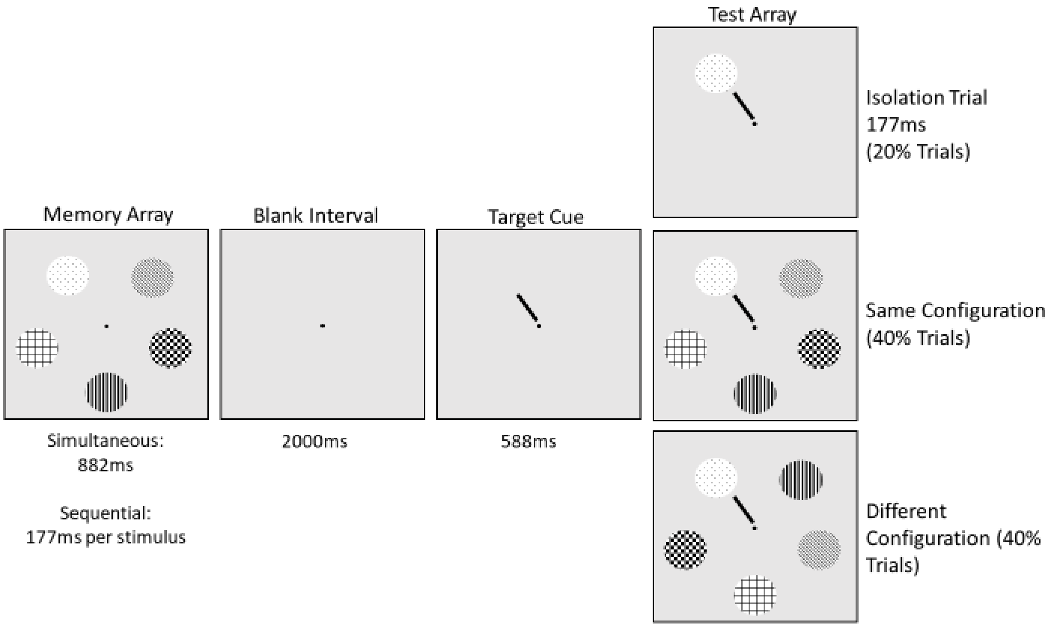

Figure 1: Trial Schematic of Overall Trial Procedure

Figure 1 shows a schematic of the general trial procedure. At the start of each trial, the fixation point would turn to a solid white, to act as a temporal alerting cue. After 706ms, the memory array would be presented. On 'sequential' trials, each item was presented for

177ms with no interval between each, while on ‘simultaneous’ trials, all items were presented together for 882ms, therefore keeping the same total array duration in both conditions. After a 2000ms blank interval, the cue would appear for 588ms, indicating at which location the target had been. Immediately following the offset of the cue, the test array would be presented. As with the memory array, if the test array was presented sequentially then each item was presented for 177ms with no interval between each, while if presented simultaneously, all items were displayed together for 882ms. On ‘isolation’ trials only the target probe itself was presented, with no other array items. To ensure similar temporal properties across conditions, isolation trials were created as sequential trials but where all the non-probe items had the same colour as the background (so invisible). This meant that the probe itself would only ever be presented for 177ms, but its exact onset would be jittered slightly, depending on its order in the sequence. After the test array offset, the fixation disc changed to a hollow ring, indicating that a response could now be made. If any response was made prior to this point, it was ignored. Participants were told to prioritise accuracy and there was no advantage to responding quickly. When a participant made an incorrect response, a small beep would sound. The next trial would commence immediately following each trial, with the 706ms delay prior to the memory onset also acting as an ITI.

At the start of each session participants first completed a practice block of twelve trials, each from a randomly chosen experimental condition. In total there were 300 trials (30 per condition), split equally into 12 blocks. Between each block participants were given the opportunity to take a break.

Design

This experiment employed a 2x2x2 repeated measure design with the factors of memory array type (simultaneous or sequential), test array type (same as memory array or different) and non-cued colour locations (same or shuffled). The isolation trials were not included in the main analyses but were still included to provide a baseline such that if any significant differences were found between configuration types, these could be compared against the isolation conditions separately to determine whether certain configuration conditions lead to enhancement or inhibition (or both) relative to the isolated probe. The dependent variables were accuracy (measured as percent correct), d’ (“d prime”) and their criterion.

Data Analysis

Accuracy scores were calculated for each condition. Participants with an average score of below 60% would have been excluded but none failed this criterion. However, one participant was excluded because their performance on at least one experimental condition was more than 2.5 standard deviations from the group mean (not considering the isolation conditions, though there were no outliers in those conditions).

Therefore, nineteen participants were included in the final analysis and data were analysed using a repeated measure analysis of variance (ANOVA) with the factors of memory array type (simultaneous or sequential), test array type (same as memory array or different) and non-cued colour locations (same or shuffled).

D' scores were calculated (for all experiments in this chapter) as d' = z(H) – z(FA), and criterion was calculated as c = -0.5 x (z(H) + z(FA)), where H is the hit-rate, FA is the false alarm rate, and z is the normalised score. Note that corrections were applied where H=1, transforming it to H=1-(1/(2*T)), and where FA=0, transforming it to FA=1/(2*T), where T is the number of trials (12).

Results

Descriptive statistics summarising the accuracy (% correct), sensitivity (d') and criterion values are given in Tables 1, 2 and 3 for each condition.

Table 1: Table displaying means (and standard deviations) of accuracy (%).

Memory Array Presentation Type	Probe Array Presentation Type	Non-cued items		
		Same	Different	Isolation
Simultaneous	Simultaneous	82.5 (8.7)	82.7 (12.3)	82.7 (9.2)
	Sequential	81.2 (10.2)	83.1 (9.1)	
Sequential	Simultaneous	74.4 (10.1)	72.0 (10.4)	73.5 (11.3)
	Sequential	72.6 (10.0)	72.0 (10.2)	

Table 2: Table displaying means (and standard deviations) of d'.

Memory Array Presentation Type	Probe Array Presentation Type	Non-cued items		
		Same	Different	Isolation
Simultaneous	Simultaneous	2.13 (0.74)	2.15 (0.98)	2.12 (0.69)
	Sequential	2.00 (0.86)	2.09 (0.76)	
Sequential	Simultaneous	1.52 (0.67)	1.30 (0.73)	1.47 (0.77)
	Sequential	1.36 (0.63)	1.30 (0.69)	

Table 3: Table displaying means (and standard deviations) of criterion values.

Memory Array Presentation Type	Probe Array Presentation Type	Non-cued items		
		Same	Different	Isolation
Simultaneous	Simultaneous	0.10 (0.44)	-0.05 (0.31)	-0.29 (0.33)
	Sequential	-0.10 (0.36)	-0.20 (0.18)	
Sequential	Simultaneous	-0.04 (0.50)	-0.05 (0.33)	-0.32 (0.36)
	Sequential	-0.16 (0.44)	-0.25 (0.29)	

Results of the statistical analysis are presented in Table 4. The only statistically reliable differences were that performance was better when the original memory array was presented simultaneously versus sequentially (both as in terms of accuracy and sensitivity), and participants adopted a more conservative criterion when the probe array was presented simultaneously. No other differences or interactions were statistically significant.

As a secondary analysis, additional 2x3 ANOVAs were also run on the data, with the factors of memory array (simultaneous or sequential) and probe array (isolation, simultaneous or sequential). The latter factor did not separate conditions for configuration change or no-change as the primary analysis found no apparent difference between these conditions. The primary purpose of this analysis was to allow for the comparison of the various probe array presentation types with the isolation condition.

For both accuracy and sensitivity, the improved performance when the memory array was simultaneous versus sequential was observed again (Accuracy: *F*(1,18)=33.700, *p*<.001, , η_p^2=.652; Sensitivity: *F*(1,18)=39.836, *p*<.001, η_p^2=.689) but, for both, no reliable differences were found between the different probe array presentation types (Accuracy: *F*(2,36)=0.188, *p*=.830, η_p^2=.010; Sensitivity: *F*(2,36)=0.563, *p*=.575, η_p^2=.030) and no interaction was found (Accuracy: *F*(2,36)=0.034, *p*=.966, η_p^2=.002; Sensitivity: *F*(2,36)=0.084, *p*=.920, η_p^2=.005). For criterion values, there was again no difference between the memory array presentation types (*F*(1,18)=1.202, *p*=.287, η_p^2=.690) but there was again a significant difference between probe array presentation types (*F*(2,36)=12.641, *p*<.001, η_p^2=.413). Pairwise comparisons repeated the previously observed more conservative criterion when the probe array was presented simultaneously than when sequentially (*p*=.001), but also compared to the isolation condition (*p*<.001), while there was no reliable difference between the sequential presentation and isolation condition (*p*=.068). There was also still no interaction between memory and probe array presentation types (*F*(1.46, 26.287)=0.097, *p*=.848, η_p^2=.005).

Table 4: results of each of the 2x2x2 ANOVAs on accuracy, sensitivity and criterion.

	Accuracy (% Correct)	**Sensitivity (d')**	**Criterion**
Memory Array Type	**F(1,18)=40.087, p<.001*,** η_p^2=.690	**F(1,18)=41.741, p<.001*,** η_p^2=.699	F(1,18)=1.525, p=.233, η_p^2=.078
Probe Array Type	F(1,18)=0.328, p=.574, η_p^2=.018	F(1,18)=1.118, p=.304, η_p^2=.058	**F(1,18)=10.492, p=.005*,** η_p^2=.368
Configuration Change	F(1,18)=0.055, p=.817, η_p^2=.003	F(1,18)=0.516, p=.482, η_p^2=.028	F(1,18)=1.676, p=.212, η_p^2=.085
Memory Array Type * Probe Array Type	F(1,18)=0.320, p=.860, η_p^2=.017	F(1,18)=0.006, p=.941, η_p^2<.001	F(1,18)=0.068, p=.797, η_p^2=.004
Memory Array Type * Configuration Change	F(1,18)=1.228, p=.282, η_p^2=.064	F(1,18)=1.364, p=.258, η_p^2=.070	F(1,18)=0.790, p=386, η_p^2=.042
Probe Array Type * Configuration Change	F(1,18)=0.767, p=.383, η_p^2=.041	F(1,18)=0.449, p=.511, η_p^2=.024	F(1,18)=0.007, p=.934, η_p^2<.001
Memory Array Type * Probe Array Type * Configuration Change	F(1,18)=0.003, p=.958, η_p^2<.001	F(1,18)=0.109, p=.745, η_p^2=.006	F(1,18)=0.351, p=.561, η_p^2=.019

Discussion

Overall results were not as predicted. The key result was that there were no effects observed regarding whether the non-cued items had the same or different configuration between the memory and test arrays. Although this may initially appear to suggest that the mechanisms underlying the configuration effect (such as ensemble encoding and relative encoding) may not operate across time as they do across space, the current study also failed to replicate these effects using even when simultaneous presentation in both memory and probe arrays. This might indicate that the presentation method may not have been the critical factor in the lack of effects observed in the sequential condition.

The lack of temporal effects contradict previous results such as Olivers & Schreij (2014), who demonstrated that the detection of a feature change in a probe array could be affected by features that temporally preceded the relevant time window. However, one key difference is that in that study, the preceding information (the direction that the array transitioned onto the screen from) affected every item in the array. It is possible that this integration of temporally distinct information was represented for each individual object independently of other objects in the group. This might suggest that the memory of a dynamically changing stimulus also incorporates information of the state of that object for a short temporal window either side of a critical period (as in Albrecht & Scholl, 2010). However, the current results might imply that information at a group level, between different items, is not similarly encoded across time.

However, as mentioned above, there is reason to be cautious about drawing any such conclusion from the current results alone as this study also found no effects of configural changes even when both the memory and probe arrays contained only simultaneously presented stimuli, as in the various paradigms demonstrating such configuration changes in previous studies (Boduroglu & Shah, 2006, 2009; Delvenne & Bruyer, 2006; Gmeindl et al., 2011; Jiang et al., 2000; Mutluturk & Boduroglu, 2014; Papenmeier et al., 2012; Sanocki et al., 2010; Silvis & Shapiro, 2014; Vidal et al., 2005; Zimmermann et al., 2010). This failure to observe even well-established configuration effects between simultaneous arrays might indicate that there are other key differences with the current design that preclude between-item relative encoding generally, and these might also have caused the lack of any effects in the sequential conditions. In order to conclude convincingly that equivalent relative encoding does not occur across temporally separate items, it would be stronger if this absence was demonstrated using a paradigm that simultaneously does yield the expected effects with simultaneous presentation.

One possible difference between the current study and those previous studies that demonstrated such an effect is the current use of a retro-cue (presented during the blank memory interval, before the test array reappears) rather than a post-cue (cueing which item is the target only as, or after, the test array is presented). Although the SOA between cue and test array was relatively short (588ms), it is possible that the advance knowledge of which item would be tested would allow for a focussing on that memory representation at a cost of others. If the memory representation were to transition from a group level to an individual object, it is possible that also any relative between-item encoding would also deteriorate and potentially no longer have a reliable influence on identification of the cued item. Alternatively, even if such information were still maintained within memory and was not 'forgotten' in such a short period of time as the above pre-cue interval, the above results might instead suggest that the utility of this information is sensitive to attentional biases. Some previous studies have suggested that integration of information may not be robust with elements that are irrelevant (Udale, Farrell, & Kent, 2017; Zhao et al., 2014), and may require that the relevant items are attended (Gronau & Shachar, 2014). Although these studies were focussed on information that was either never task relevant, or involved semantic associations, similar effects could apply here, such that the removal of task relevance, and any associated attentional resources, from the non-cued items may have released or removed any previously integrated information.

Even if ensemble information is still maintained even when spatial attention is focussed on just a single item, its manifestation might differ from when attention is spread diffusely across multiple items. It might therefore be the mismatch of attentional deployment between encoding and test phases (diffuse attention across multiple items during the encoding array but focussed attention during the test array due to the use of the retro-cue), rather than strictly the sequential or simultaneous presentation types, that limit the effective utility of any encoded information in the current study. Overall, it might be that such information simply being encoded to memory is not sufficient for its effects to manifest and that any such influence on memory performance may depend on the attentional state. If attention is focussed towards an individual object within memory then the between-object information may be unattended or even inhibited, therefore yielding reduced influence on behaviour than when attention is still more diffuse and spread among the entire array.

This may seem surprising as ensemble encoding appears to operate efficiently under conditions of reduced attention (Alvarez & Oliva, 2008, 2009). However, in those designs, responses were still occasionally required to the less attended stimuli, whereas our cue had 100% validity and the other items were never required for response. Such mechanisms may therefore still operate even on task-relevant items with reduced attention but not across items of no task relevance. Alternatively, as outlined previously, the exact mechanisms underlying the configuration effect are unclear and it may be that mechanisms such as computing the average value of an array, even if they do operate under conditions of reduced attention, are not alone sufficient to exhibit the predicted behavioural differences. For example, the current results might indicate that some of the parallel mechanisms, such as relative encoding between items, do not operate efficiently between task-irrelevant items.

Either way, if the timing of the cue is a critical factor in the difference in results between this study and previous ones, it remains possible that even with sequentially presented objects, if the cue were not presented in advance of the memory array, that between-object encoding might also have been evident. To test this, the following studies will include both a retro and post cue. Both types of cue will be included to test explicitly whether the timing of the cue does influence configuration effects in memory.

An alternative issue with the current paradigm might be that of the complexity and ‘mixed’ nature of its conditions. Within each block every trial would randomly be drawn from one of 10 different experimental conditions, not considering the different target locations. The sheer variety and unpredictability of each array being simultaneous or sequential might have hindered effective encoding strategies. Although overall performance seemed reasonable, such that the lack of effects might not simply reflect the task having been too difficult, it could nevertheless be the case that between-item relative encoding is not truly automatic and instead is dependent on the situation. For example, in the current experiment, over half of the trials had non-simultaneous probe array presentations. If such between-item relative encoding is only effective for items presented simultaneously then this might mean there would have been only a very limited advantage to any additional effort to form such representations during the encoding period, whereas in previous studies using only simultaneous presentation, it would have been certain (aside from the minority of ‘isolation’ trials) that any encoded between-object information would be able to be utilised during a simultaneously presented probe array. A similar possibility exists even if such between-item encoding is possible also for sequentially presented stimuli. If any such between-object representations have some specificity for the type of presentation, so spatial encoding of simultaneously presented stimuli does not generalise to sequentially presented stimuli and vice versa, then again there may have been limited advantage to any such encoding in the current experiment as such representations could be utilised on less than half of trials. There are previous examples of apparent configuration effects where the encoding array was presented sequentially and the test array always presented simultaneously (Blalock & Clegg, 2010), but as outlined earlier, there remain other confounding factors that could account for their results. Previous studies showing spatiotemporal configuration-like effects used designs where each block contained only a single presentation type (Ihssen et al., 2010; Silvis & Shapiro, 2014). To test this, the next studies will also therefore employ a ‘blocked’ design, such that all trials in each block use either a simultaneous or a sequential presentation type consistently.

Although the main result of this study was the lack of any configuration change effects, there are also other secondary results that require discussion. Perhaps the most obvious is that performance was reliably better (whether measured as accuracy or sensitivity) when the memory array was presented simultaneously relative to when it was presented sequentially. Previous studies have sometimes found memory performance to be better when using sequential displays than simultaneous displays (Ahmad et al., 2017; Ihssen et al., 2010), but only when the sequential presentation allows sufficient consolidation time and is presented in a predictable manner. Crucially, any stimuli presented sequentially were done so in a predictable order at predictable locations. In designs where the order and locations of sequential stimuli are unpredictable, as in the current study, performance tends to be worse for sequential presentation than simultaneous (Blalock & Clegg, 2010; Lecerf & De Ribaupierre, 2005). Furthermore, those finding a sequential advantage often use slower or longer presentation times, allowing for greater consolidation of information (Ricker & Cowan, 2015), whereas the rate of stimulus presentation in the current study was relatively fast.

The current finding of a simultaneous presentation advantage may therefore not be a surprising result as it is likely that encoding a simultaneous array may be easier than a sequential array given the current stimulus parameters. Although the two arrays were matched so that their total presentation time was equivalent, this would only apply if encoding even of the simultaneous array were fully serial (such that each item was encoded sequentially). If any encoding occurred in parallel, then the additional time that each individual item was present for would have enabled better encoding. Even if there was a serial component to encoding, the simultaneous presentation would also have allowed for more efficient encoding due to more predictable shifts in spatial attention – all items would have been available for encoding at all positions so participants could have shifted their attention between them in whatever manner they wished. However, in the sequential presentation, each item appeared in a random location at a relatively fast rate and if participants attempted to shift their attention between each item as it appeared then there would have been a relative cost in having to detect the location of the new stimulus, to shift spatial attention to that location and then attempt to encode its features in the time remaining before it vanished and the next stimulus appeared. A further possibility might be that in the sequential condition, there may be additional information to encode if any between-item encoding does take place, such as the temporal order, at a cost to other information encoding and therefore leading to poorer performance. However, if such additional information truly were encoded then it remains unclear why it would not have led to configuration-based differences in terms of the probe array presentation types.

The other significant result of the current study was that participants adopted a more conservative criterion (more likely to report that the target had not changed) when the probe array was a simultaneously presented array rather than a sequential array or isolated item. This effect did not appear to depend on the configuration of that array, nor on the initial presentation of the memory array, and did not seem to affect overall performance (in terms of accuracy or sensitivity). The reason for this pattern is not immediately obvious based on existing research.

The next two experiments will attempt to address the issues identified by using only simultaneous (experiment 2) or sequential (experiment 3) presentation types, to investigate whether a more simplified design may encourage the encoding and utility of more reliable group-level information. It should be noted that this need not be purposeful encoding strategies but could also be due to implicit statistical learning. Both of the next experiments will also include both a retro-cue and post-cue condition to test explicitly whether differences arise from the ability to focus attention on an individual stimulus relative to still needing to maintain multiple items at the moment of the onset of the test array.

Experiment 2: Attempt to Replicate Configuration Effects with Purely Simultaneously Presented Arrays

Introduction

The next experiment was designed to build on the results of the previous experiment to investigate whether the timing of the cue relative to the test array, and the predictability of presentation types, might affect the utilisation of configural information in memory for groups of items.

In order to address the issue of predictability, the current experiment will use only simultaneous presentation, never sequential, to investigate whether the timing of the cue affects the previously established spatial configuration effects (Jiang et al., 2000; Vidal et al., 2005). This is to attempt to replicate previous configuration results using similar simultaneous arrays as have successfully yielded configuration results previously, while manipulating only cue timing. If such results are replicated the intention would then be to run a similar study using only simultaneous presentations for memory and probe arrays.

To address the issue of cue timing, the following study will have two cue conditions: one will use a retro-cue, identical to the previous study, while the other will employ a post-cue. A post-cue is typically a cue whose onset is either simultaneous with, or after, the presentation of the test array. In this study, the post-cue will appear simultaneously with the rest of the probe array, similar to previous studies showing configuration based effects (Vidal et al., 2005).

Based on previous studies employing different cue types in working memory paradigms (independent of configuration or ensemble encoding effects), it is likely that overall performance will be higher on retro-cue trials (Souza & Oberauer, 2016) as these allow the target to be known slightly earlier, reducing the demand to maintain irrelevant non-target items for as long. It should be noted that the retro-cue advantage is not simply due to a shorter retention period and shows an advantage even if the time between the encoding array and retro-cue is the same as the time between the encoding array and test array/post-cue (such that the overall retention period is longer in the retro-cue condition than post-cue; Pinto et al., 2016). The critical feature seems for allowing time for focussed attention to consolidate further the target representation and perhaps also protect it from interference from the presentation of the test array.

It remains possible that, despite the lack of configuration-based effects observed in the previous experiment, they may still be observed here even in the retro-cue condition. This is because, if the complex mixed design in the previous study was a major cause of not observing any configuration-based changes, then the simpler blocked design of this study might allow for any such effects to be observed, though even if present, there may still be a relative difference in performance between the cue types.

It is predicted that memory performance will be greater when the non-cued items have the same spatial configuration versus when they have a different configuration, with performance on trials with only an isolated probe stimulus (no other array items) somewhere between the two. However, based on the results of the first study, it is predicted that these differences will be greater on trials using a post-cue and not retro-cue trials (and may even be absent entirely). It is also expected that overall performance will be higher on retro-cue than post-cue trials.

Method

Participants

Twenty-one participants (5 males, 16 females) aged between 18 and 30 (M=20.4, SD=2.6) volunteered in return for either course credit or £5 (for 30 minutes). All were right-handed and reported normal (or corrected to normal) vision. All procedures were reviewed and approved by the Departmental Ethics Committee.

Stimuli and Apparatus

The stimuli and apparatus were identical to the previous experiment.

Procedure

The procedure was identical to that of the previous experiment, with only three exceptions.

In the current study, stimuli in both the memory and probe arrays were always presented simultaneously, never sequentially. Critically, the cue could now appear either 588ms prior to the onset of the probe array (retro-cue), as in the previous experiment, or at the same time as the probe array (post-cue). The cue duration (588ms) was the same in both conditions, such that a post-cue was not displayed for longer than a retro-cue, even if the probe array remained visible for longer until the participant made a response.

Also, now one third of trials were 'isolation' trials (where the probe stimulus was presented alone with no non-cued items) so that there was an even probability of each trial being one of the three possible probe array types (same configuration, different configuration, and isolation). Note that other parameters (such as pre or post cue, target position and target change) were still counterbalanced within the isolation trials.

In total there were 240 trials (40 per experimental condition), split equally between eight blocks. However, the trial conditions in the first four blocks and the latter four blocks were calculated and counterbalanced independently of each other. Between each block participants were given the opportunity to take a break.

Design

This experiment employed a 2x3 repeated measure design with the factors of Cue Type (retro-cue or post-cue) and Probe Array Configuration Type (same configuration, different configuration, or isolation). As in the previous experiment, the dependent variables were accuracy of responses (%), d' and criterion on the change detection task.

Data Analysis

Accuracy scores were calculated for each condition. Any participant whose average score was below 60% would have been excluded from further analyses but no participants met this condition. Any participants with performance more than 2.5 standard deviations from the group mean in any of the individual conditions were excluded (2 participants).

Therefore, nineteen participants were included in the final analysis and data were analysed using a 2x3 repeated measure analysis of variance (ANOVA) with the factors of Cue Type (retro-cue or post-cue) and Probe Array Configuration Type (same configuration, different configuration or isolation).

Results

Table 5 illustrates the accuracy results:

Table 5: Mean (Standard Deviation) Descriptive Statistics for Accuracy (% Correct)

		Cue Type		
		Retro-cue	Post-cue	Total
Probe Array Configuration Type	Isolation	87.4 (6.9)	82.5 (9.6)	84.9 (7.1)
	Same	90.79 (5.90)	84.3 (8.37)	87.6 (6.4)
	Different	88.16 (6.8)	81.2 (9.0)	84.7 (6.8)
	Total	88.8 (5.2)	82.7 (7.4)	

There was no significant main effect of probe array configuration (though note it could be considered marginal; *F*(2,36)=2.996, *p*=.063, η_p^2=.143) and no interaction between probe array configuration and cue type (*F*(2,36)=0.375, *p*=.690, η_p^2=.020). However, performance was significantly higher with retro-cues than with post-cues *F*(1,18)=31.445, *p*<.001, η_p^2=.636).

In addition to percent correct, signal detection calculations were also performed for d' (Table 6) and criterion (Table 7).

There was no significant effect of configuration (though again note potentially marginal; *F*(2,36)=2.893, *p*=.068, η_p^2=.138) and no interaction between cue type and configuration (*F*(2,36)=0.863, *p*=.430, η_p^2=.046). However, d' was significantly higher when a retro cue was used than when a post-cue was used (*F*(1,18)=32.836, *p*<.001, η_p^2=.646).

Table 6: Mean (Standard Deviation) Descriptive Statistics for d'

		Cue Type		
		Retro-cue	Post-cue	Total
Probe Array Configuration Type	Isolation	2.52 (0.63)	2.21 (0.71)	2.37 (0.59)
	Same	2.83 (0.64)	2.26 (0.77)	2.54 (0.64)
	Different	2.57 (0.70)	1.98 (0.77)	2.27 (0.61)
	Total	2.64 (0.54)	2.15 (0.60)	

Table 7: Mean (Standard Deviation) Descriptive Statistics for Criterion Values

		Cue Type		
		Retro-cue	Post-cue	Total
Probe Array Configuration Type	Isolation	-0.33 (0.20)	-0.52 (0.20)	-0.43 (0.15)
	Same	0.00 (0.31)	-0.17 (0.34)	-0.08 (0.22)
	Different	-0.17 (0.28)	-0.29 (0.25)	-0.23 (0.15)
	Total	-0.17 (0.17)	-0.32 (0.18)	

It is worth noting that the average criterion scores were negative (liberal) in almost every condition (except for same configuration retro-cue, which had an average criterion value of 0.00). There was a significant main effect of configuration ($F(1.481, 26.654)=18.978$, $p<.001$**,** $\eta_p^2=.513$), with criterions significantly more conservative when the non-cued items had the same configuration than when absent entirely ($p<.001$) or when with a different configuration ($p=.023$). Criterions were also significantly more conservative when non cued items had a different configuration than when absent ($p=.002$). Criterions were also significantly more conservative when a retro cue was used than when a post cue was used (*F*(1,18)=6.312, *p*=.022, η_p^2=.260), but there was no interaction between configuration and cue type (*F*(2,36)=0.222, *p*=.802, η_p^2=.012).

Discussion

Overall the results of this study still failed to replicate any effects from changes in configuration of non-target items, despite using only simultaneous presentation of arrays, as primarily used in previous studies (Jiang et al., 2000; Vidal et al., 2005). However, the differences between configuration conditions, in both accuracy and d prime, might be considered 'marginal'. It must be emphasised that ultimately this result should not be misinterpreted as sufficient evidence to conclude that performance improved when non-targets had the same rather than different configurations, but neither does the evidence suggest a total absence of any such effects and further research will be required to confirm whether this pattern can be replicated more reliably or is shown to be a statistical anomaly.

Regardless of this ambiguity, there was no evidence of any interaction with cue type. This means that even if the previous result might indicate weak configuration effects, this potential difference was not unique to only one of the two cue types. Taken together, these results indicate that the lack of configural effects in the previous study was unlikely to have been due to the use of a retro cue in that experiment. This study has therefore found no configural effects even while using a post-cue, and so other explanations would be needed for the lack of configuration differences in both this and the previous result. However, once any other factors precluding configuration effects are addressed, differences between cue types might then be observed.

Given that the current evidence does not suggest that cue type may have been the major reason for a lack of any configuration differences in the previous experiment, the other suggested explanation was that only a minority of trials in experiment 1 had the same presentation type at both memory and test, which may have discouraged the encoding or utilisation of such group-level representations. However, this explanation could not account for the lack of predicted differences in this experiment as all trials used consistently simultaneous presentation.

Neither of the proposed reasons for the lack of expected results in experiment 1 (the use of a retro-cue and a complex mixed design) may account for the lack of such effects in the current study, suggesting that alternative issues may be the cause. It is not clear what these other differences might be though as the current study did not differ markedly from previous studies that did show such effects (Vidal et al., 2005) except for the inclusion of a retro-cue condition and mixed design.

It remains possible though that the inclusion of the retro-cue condition may have weakened any such effects even for the post-cue trials. As already outlined, if ensemble information is only reliably utilised when multiple items are still being maintained in memory at the moment the test array is presented and the comparison is made then on the half of trials that use a retro-cue, and allow for the focussing on just a single item, summary statistics may have limited utility. Furthermore though, on trials that use a post-cue and do continue to maintain multiple items at the moment of the comparison with the test array, any maintained summary statistics would not be helpful on trials where the test array had a changed configuration anyway, or when the probe stimulus was presented in isolation. Therefore, ensemble encoding may have a positive use on only 1/6 trials, which may discourage such encoding, either purposefully or implicitly. Future research may wish to take the blocked design approach even further such that all trials in each block use only either a retro-cue or post-cue, not a mix of both, in order to allow for maximal presentation predictability. Even if such an approach were to yield more reliable configuration effects though, the current studies would imply that such effects may not be automatic but depend on various parameters and strategies. The next study will still mix retro-cue and post-cue trials within the same block though to allow for more reliable comparison across experiments.

The other main result was that performance was better overall on trials using a retro-cue than those using a post-cue, as was predicted. This is in line with previous research (Souza & Oberauer, 2016) and could simply be due either due to reduced memory demand on retro-cue trials, or being better able to protect the key representations from interference from the onset of the test array. The retro-cue identifies which location is the target slightly earlier than the post-cue, reducing the length of time during which participants need to maintain the other non-target items in memory and therefore lessening the chances of memory degradation or interference. It should be noted though that studies focussed on measuring the retro-cue advantage explicitly will often account for this by equating the total interval between the encoding array and the retro-cue, with the duration between the encoding array and test array (and post-cue), by making the overall delay period shorter in the latter condition (Pinto et al., 2016). This was not done in the current study as any retro-cue advantage was not the focus. Aside from the time the full array needs to be maintained, the retro-cue can also help to protect key representations from disruption when the probe array itself is presented. The retro-cue means that participants have advance notice to focus their attention and need to maintain and protect only one item from such interference, whereas the post-cue requires all five items to be maintained and protected from the onset of the probe array. The latter is more difficult and so more likely to result in increased interference and therefore less accurate change detection performance.

Although the above results have failed to find any reliable effects on memory performance due to differences in probe array configuration, the current study did show that configuration differences could affect some aspects of participants' behaviour and decisions. Participants were significantly more conservative in their responses (more prone to responding that a change did not occur) when the non-cued items were in the same configuration than when either different or absent. This is in line with previous studies (Boduroglu & Shah, 2009) and could be due to the number of 'changes' within the array between each condition. When the non-cued items change configuration then there are four array changes (or five if the target also changes) between memory and probe array, while in the same-configuration condition there are only one or no changes in the array. Therefore, in the 'change' condition, the additional impression of change from the four changed non-target items might bias participants to report a 'change' (more liberal criterion) relative to when the non-cued items all did not change. Alternatively, even if not the case that one of the non-cued items was explicitly confused with the target, the noise from the additional changes across the array could also increase the apparent evidence and make it more likely that the threshold for giving a 'change' response is met.

Although the current study did not replicate traditional configural effects using simultaneous presentation, to complete the investigation of these effects, the next experiment used a similar design but with sequential presentation of arrays

Experiment 3: Configuration Effects When Using Only Arrays of Sequentially Presented Stimuli

Introduction

The previous two experiments have failed to find any effects on memory from differences in the configuration of non-cued items on memory performance for the target items. Although in the first study this might have been accounted for by issues in the design, such as the use of a retro-cue and the complexity that arises from the inclusion of so many conditions, the second experiment failed to demonstrate the predicted effects using a similar design as previously published research. This suggests that the lack of configuration effects in the first study was not simply due to the use of retro-cues rather than post-cues.

The current study will use an almost identical design to that of experiment 2 but will instead use only sequential presentation for all arrays. This therefore still addresses the identified issues from experiment 1, such as the use of only retro-cues and the complex mixed blocks design. However, the use of sequential presentation may also lead to different results than the simultaneous presentations of experiment 2; especially relating to any issues from the mismatch of attentional focus between encoding and test arrays on the majority of trials, and any negative effects that may have had on the utilisation of ensemble information.

As outlined earlier, it may not be the mismatch between presentation types that was the main issue in experiment 1 (though how well encoded information generalises between presentation types remains an interesting question, but which is outside the scope of the current experiment), but also perhaps the relative mismatch in attentional focus during the two critical periods. Any encoded ensemble information may manifest differently when attention is diffuse across multiple spatially separate items than when focussed on a single item. It might be therefore that differences in such information might limit or prevent its reliable utility when the attentional focus changes between encoding and test. In experiment 2 this occurred on over half of the trials (all the trials where a retro-cue allowed attention to focus on just a single item in anticipation of the probe array, plus any post-cue ‘isolation’ trials) and of those remaining trials where attention was spread across multiple items during both encoding and test arrays, only half of those (1/6th of total trials) had the same configuration, and so allowing ensemble information to be useful. This potential lack of trials in which ensemble information could have been useful may have reduced efforts (explicit or implicit) to rely on them.

However, when using sequential presentations, as in the current study, this may be less of an issue as there will always be only one item visible on screen at any time, with the full focus of attention on it (except for any attentional resources committed to maintaining already presented items in memory). This would lead to far less of a mismatch between attentional deployment between encoding and probe arrays compared to experiment 2. There still would be important differences between retro-cue and post-cue trials, where the latter still requires the maintenance of all five items while the former allows attention to focus on only one, and these may again weaken any predicted results. However, the ability always to focus on only one item in a given moment regardless of array or trial type, means that these differences are likely to be far less than they were in experiment 2, where most of the encoding arrays had spatially diffuse attention across multiple simultaneous items while most of the test arrays had attention focussed on only a single item. It might therefore be expected that configuration effects could be stronger when using sequential displays in the current experiment.

It is predicted that performance will be better when the non-cued items keep the same configuration between memory and probe arrays than when they change configuration (with isolation conditions falling somewhere between the two). Based on the results of the previous experiment, this difference is not predicted to differ between the two cue conditions, though it is expected that overall performance will be better on trials with a retro-cue than those with a post-cue.

Method

Participants

Twenty participants (8 males, 12 females) aged between 18 and 25 (M=20.3, SD=2.3) volunteered in return for either course credit or £5 (for 30 minutes). All were right-handed and reported normal (or corrected to normal) vision. All procedures were reviewed and approved by the Departmental Ethics Committee.

Stimuli and Apparatus

The stimuli and apparatus were identical to the previous experiment.

Procedure

The procedure was identical to that of the previous experiment, with the only exception that all memory and probe arrays were sequential rather than simultaneous, as in the sequential conditions of Experiment 1. As in Experiment 1, the configuration of the non-cued items applied both to the spatial positions of each colour but also the temporal order in which they were presented. As in Experiment 2, when this configuration was the same, every non-target item kept identical parameters in both the memory and probe arrays, and when the configuration changed, no non-target item had the same spatial or temporal value in both the memory and probe arrays.

The non-cued items were again blanked by making them the same colour as the background, but the parameters otherwise remain the same. This means that in the isolation condition, the target might not have appeared at the same time as the post-cue if the target was one of the later items in the sequence, and so even in the post-cue condition, the cue may still have appeared as similar to the retro-cue. A total of 240 trials were run (40 per condition).

Design

This experiment employed a 2x3 repeated measure design with the factors of Cue Type (retro-cue or post-cue) and Probe Array Configuration Type (same configuration, different configuration, or isolation). The dependent variables were accuracy of responses (%), d' and criterion on the change detection task.

Data Analysis

Accuracy scores were calculated for each condition. Any participant whose average score was below 60% was excluded from further analyses (4 participants). Any participants with performance more than 2.5 standard deviations from the group mean in any of the individual conditions would also have been excluded but no participants met this condition.

Therefore, sixteen participants were included in the final analysis and data were analysed using a 2x3 repeated measure analysis of variance (ANOVA).

Results

Table 8 displays the mean (and standard deviation) performance levels for each condition. There was a significant main effect of configuration type ($F(2,30)=3.683$, $p=.037$, $\eta_p^2=.197$). This effect was driven by accuracy being higher in the isolation condition than in both the 'same' and the 'different' configuration conditions. However, neither of these post hoc comparisons survived Bonferroni corrections for multiple comparisons ($p=.078$ for the 'same' configuration condition, and $p=.082$ for the 'different' configuration condition, with 'same' and 'different' conditions not different from each other, $p>.999$). Performance was also significantly better when a retro-cue was used than when a post-cue was used ($F(1,15)=10.799$, $p=.005$, $\eta_p^2=.419$) but there was no interaction between configuration type and cue type ($F(2,30)=0.460$, $p=.636$, $\eta_p^2=.030$).

Table 8: Mean (Standard Deviation) Accuracy (% Correct) Descriptive Statistics

		Cue Type		
		Retro-cue	Post-cue	Total
Probe Array Configuration Type	Isolation	79.22 (9.9)	75.63 (10.3)	77.42 (9.4)
	Same	76.88 (9.24)	71.25 (11.11)	74.06 (9.3)
	Different	76.88 (9.24)	70.78 (10.79)	73.8 (8.6)
	Total	77.66 (8.3)	72.55 (9.6)	

In addition to percent correct, signal detection calculations were also performed for d' (Table 9) and criterion (Table 10).

Table 9: Mean (Standard Deviation) descriptive statistics for d':

		Cue Type		
		Retro-cue	Post-cue	Total
Probe Array Configuration Type	Isolation	1.91 (0.81)	1.56 (0.77)	1.73 (0.74)
	Same	1.62 (0.72)	1.27 (0.81)	1.45 (0.71)
	Different	1.60 (0.69)	1.18 (0.63)	1.39 (0.56)
	Total	1.71 (0.67)	1.34 (0.65)	

In line with the accuracy results, there was a significant main effect of configuration (*F*(2,30)=5.878, *p*=.007, η_p^2=.282), driven by the isolation condition having significantly better performance than the same configuration (*p*=.018) and the different configuration (*p*=.010), though the latter two were not different from each other (*p*>.999). Performance was also significantly better with the retro cue than the post cue (*F*(1,15)=13.779, *p*=.002, η_p^2=.479).

However, there was no significant interaction between the cue type and configuration type (*F*(2,30)=0.102, *p*=.904, η_p^2=.007).

Table 10: Mean (Standard Deviation) Descriptive Statistics for Criterion Values:

		Cue Type		
		Retro-cue	Post-cue	Total
Probe Array Configuration Type	Isolation	-0.23 (0.43)	-0.18 (0.30)	-0.21 (0.33)
	Same	-0.00 (0.34)	-0.07 (0.36)	-0.37 (0.33)
	Different	-0.20 (0.19)	-0.24 (0.18)	-0.21 (0.16)
	Total	-0.15 (0.23)	-0.16 (0.18)	

There was no significant effect of configuration type (*F*(2,30)=2.378, *p*=.110, η_p^2=.137) and no significant effect of cue type (*F*(1,15)=0.162, *p*=.693, η_p^2=.011). There was also no significant interaction between configuration type and cue type (*F*(2,30)=0.876, *p*=.427, η_p^2=.055).

Combined Analysis

As a secondary analysis, this experiment and the previous experiment were combined and subjected to a 3x2x2 mixed ANOVA with the within-subject factors of probe type (same-configuration, different-configuration and isolation) and cue type (retro-cue and post-cue), and the between-subject factor of experiment (Experiment 2 and Experiment 3).

Full statistical results are presented in Table 11. The key results are that overall, both accuracy and sensitivity were significantly better in the simultaneous-presentation experiment than the sequential-presentation experiment and significantly better on retro-cue than post-cue trials. Both measures of performance showed an interaction between experiment and probe array presentation type. This was driven by the previously described results, that in experiment 2 there was no statistically reliable difference between the probe array presentation types, while in experiment 3 the isolation condition had significantly better performance (both in percent correct and d’) than either of the full array conditions. Sensitivity was overall better in the isolation condition than in the ‘different’ configuration condition ($p=.04$), but not than in the 'same' configuration condition ($p>.999$), and there was no significant difference between the 'same' and 'different' conditions ($p=.104$), though this was not reflected in accuracy performance. Finally, criterion values were significantly more conservative on retro-cue trials than post-cue trials. No other results or interactions reached significance.

Table 11: Statistical Results of the Various Analyses Across Experiments 2 and 3 Combined.

	Accuracy (% Correct)	**Sensitivity (d')**	**Criterion**
Probe	$F(2,66)=2.165$, $p=.123$, $\eta_p^2=.062$	**$F(2,66)=4.054$, $p=.022*$,** $\eta_p^2=.109$	$F(2,66)=12.385$, $p=.118$, $\eta_p^2=.273$
Cue	**$F(1,33)=36.572$, $p<.001*$,** $\eta_p^2=.526$	**$F(1,33)=43.355$, $p<.001*$,** $\eta_p^2=.568$	**$F(1,33)=4.946$, $p=.033*$,** $\eta_p^2=.130$
Experiment	**$F(1,33)=19.01$, $p<.001*$,** $\eta_p^2=.366$	**$F(1,33)=19.39$, $p<.001*$,** $\eta_p^2=.370$	$F(1,33)=3.365$, $p=.076$, $\eta_p^2=.093$
Probe * Experiment	**$F(2,66)=4.646$, $p=.013*$,** $\eta_p^2=.123$	**$F(2,66)=4.295$, $p=.018*$,** $\eta_p^2=.115$	$F(2,66)=2.205$, $p=.118$, $\eta_p^2=.063$
Cue * Experiment	$F(1,33)=0.287$, $p=.596$, $\eta_p^2=.009$	$F(1,33)=0.798$, $p=.378$, $\eta_p^2=.024$	$F(1,33)=3.229$, $p=.082$, $\eta_p^2=.089$
Probe * Cue * Experiment	$F(2,66)=0.009$, $p=.991$, $\eta_p^2<.001$	$F(2,66)=0.398$, $p=.516$, $\eta_p^2=.012$	$F(2,66)=0.660$, $p=.520$, $\eta_p^2=.020$

Discussion

As with the previous two experiments, the results of the current experiment did not confirm the predicted patterns. Overall it was found that performance was best when the probe item was presented in isolation rather than with other items in the array, but whether those other items had the same or different configuration as in the memory array seemed to make no difference. Unlike the previous studies, there seemed to be no differences in criterion values between any conditions. Finally, replicating the results of the previous study, performance was overall better on trials using a retro-cue than those using a post-cue.

The key finding of the current study was the lack of any difference in performance based on the configuration of the probe array. Although the main effect was significant, this was due to performance being better when the probe item was presented in isolation than when other array items were presented, regardless of their configuration. This is therefore not due to configural or ensemble encoding effects but may instead reflect a difference in task difficulty between the two conditions. The presentation of a test array in a working memory paradigm always has the potential to interfere with existing memory representations. As noted earlier, part of the explanation for the retro cue advantage, independent of the absolute retention period duration, is the improved ability for attention to protect the cued representation against this interference (Souza & Oberauer, 2016). However, the current presentation type of a dynamic scene with constant onsets of new objects is perhaps especially likely to disrupt attention. Abrupt object onsets seem particularly efficient at capturing attention (Cole, Kentridge, & Heywood, 2004; Irwin, Colcombe, Kramer, & Hahn, 2000; Yantis & Jonides, 1988), even compared to other salient changes to otherwise existing objects. This continued capture of attention might have drawn attention away from the cued target location, and perhaps even from attempts at maintaining the target's internal representation, ultimately making comparisons more difficult (and therefore less accurate) than in the isolation condition where only the single cued probe is presented. Future research that intends to use sequential presentations for test arrays may benefit from the use of placeholders that remain visible throughout a trial and change properties rather than absolutely appear and disappear. The potential issue with this approach though, and the reason such an approach was not adopted in the current experiments, was to minimise the influence from such additional visible information on ensemble encoding mechanisms.

Overall performance was worse in this experiment (sequential presentation) than experiment 2 (simultaneous presentation), as reflected in the lower average performance measures and higher exclusion rate (and as confirmed statistically in the combined analysis). This replicates the results of experiment 1 that also found performance worse when the memory array was presented sequentially than simultaneously. As discussed previously, this is perhaps likely due to a combination of increased difficulty at both the encoding and retrieval stage. Sequential stimuli with the parameters used in these experiments (rapid presentation and unpredictable locations) would be unable to utilise any parallel processes and even any serial processes would have reduced efficiency by needing to detect and reorient to the unpredictable onsets of each new stimulus. In the simultaneous presentation, any such serial shifts in attention could be controlled in a top-down manner, and resources could be deployed flexibly across items according to which would more benefit from additional encoding.

Regardless, the key result is that, despite the better performance on isolation trials, there were no apparent effects of the configuration of the non-cued items when they were presented. This replicates the lack of configuration effects in experiment 2 and would suggest that issues such as the mixed trial types or the use of a retro-cue were not the cause of the lack of differences in that experiment. Instead, it seems that there may simply be no equivalent in sequentially presented arrays to the between-item encoding that has previously been observed in simultaneously-presented arrays.

Performance was once again better on trials with a retro-cue than post-cue, replicating results from experiment 2. Again, this is most likely due to reduced memory load both from the shorter period it is required to maintain the non-target items for, and for the advantage of being able to protect the key representation better prior to the onset of the test array. However, as in the previous experiment, the type of cue did not interact with probe array presentation type, again suggesting that the results of experiment 1 were not influenced by the use of a retro-cue.

Contrary to experiment 2, there were no differences between conditions observed in the criterion analyses. Although not the primary focus of this series of research, this difference in results is still relevant. In experiment 2 (simultaneous presentation only), participants' responses were more liberal when the non-cued items had a different (vs. same) configuration, replicating previous findings (Boduroglu & Shah, 2009) and likely due to interference or confusion from the 'changes' of the non-cued items. However, no such differences were observed here using sequential presentation. This might indicate that the changes of the non-cued items did not have the same level of confusion or interference as with simultaneously presented stimuli. This may therefore add further evidence that between-item representations are not formed across time between items that are temporally separate. However, the interaction was not significant in the combined analysis, so this may not be a reliable difference.

General Discussion

Together, the three experiments in this chapter aimed to test whether group level or between-item information is encoded also across temporally separate, sequentially presented items, or only between items that are presented simultaneously to each other. Overall, the research failed to find any evidence of such representations being developed in the case of sequential presentation. The only reliable difference observed between simultaneous and sequential presentation types regarding configuration was that with sequential presentation, performance was better when the test probe was presented in isolation than as part of a full array, regardless of whether the array's configuration had changed or not. This is likely just due to the more disruptive nature of continued (and potentially unpredictable) onsets and offsets of stimuli during the probe array in the sequential condition, making it harder to direct attention as effectively towards the probed item. Overall, the obvious main conclusion is therefore that between-item associations do not form across sequentially separate objects.

However, it is vital to emphasise that the studies in this chapter also failed to replicate any reliable configuration effects even on trials using only simultaneous presentation types. This makes any firm conclusion difficult as it remains possible that the reason for failing to observe between-item encoding in the sequential condition was not actually due to the sequential nature of presentation itself but instead due to secondary factors that prevented reliable between-item encoding at all, including between simultaneously presented stimuli. More convincing evidence of a lack of such effects for temporally distinct items specifically would have been to replicate the previously reported spatial effects with simultaneously presented stimuli but not when those same stimuli were presented sequentially. Given only the nature of presentation would have been different between the two conditions, it would have allowed for a confident attribution of the difference in results to the different presentation types.

However, the lack of reliable replication of previously reported spatial effects does not negate any further interpretation of the sequential results. It was never an absolute prerequisite that the conditions ideal for spatial between-item encoding would be the same as those for temporal between-item encoding, and it could have been possible that these studies could have revealed evidence of between item encoding in the sequential conditions and not in the simultaneous conditions. The two effects need not necessarily occur in parallel and some dissociation or distinction between them would not be unreasonable.

This is especially true when considering that there are two distinct mechanisms by which configuration information could operate across time. The first is the spatial between-item encoding, with the novel question being whether these spatial relationships are encoded even when each item is in isolation during the moment of its presentation. However, the second way configural effects may have manifested across sequential stimuli would be in the temporal dynamics or temporal order of changing information, regardless of the spatial position of the information. For example, it might be possible to adapt the configuration paradigm such that stimuli are presented sequentially at a single location, with the critical factor between the same and different configuration conditions being not whether the target and non-targets maintain a consistent spatial configuration, but whether they maintain the same temporal configuration (order). For example, if change detection in the memory task is affected by whether the stimuli presented immediately prior and/or after the critical stimulus keep consistent identities or not. An initial consideration for Experiment 1 had been to vary the spatial and temporal configuration conditions independently but the number of conditions and trials grew too large so the approach described was adopted where sequential stimuli had either both the same spatial and temporal properties or different spatial and temporal properties. This was to maximise the difference between the two conditions and increase the chances of detected any differences, which could then be explored further in more subtle or focussed experiments. The current studies failed to find any difference due to such posited mechanism. The associated failure to find configuration effects using simultaneous presentation is most relevant to the former mechanism, as it remains possible that parameters in these experiments precluded any spatial configuration encoding regardless of presentation type. However, these lack of effects in experiment 2 do not account so easily for the lack of any sequential effects that could have been caused by the second mechanism.

Therefore the null results of the sequential presentation conditions will still be interpreted here in contrast to previous studies, including those finding configuration effects using simultaneous presentation conditions, and possible explanations for why such effects might not have been observed here, but it is important to make clear that the current evidence cannot rule out the possibility that the null results in both presentation conditions is not due to some secondary factor and, if addressed in future research, that configuration effects would be observed in both simultaneous and sequential presentation conditions.

The current research differed in design from previous studies by employing a retro-cue, presented during the memory interval before the onset of the probe array, in contrast to the post-cues used in previous research (Vidal et al., 2005), which are presented after the memory interval, at the onset of the probe array. It was initially thought that this use of retro-cues may have affected the pattern of results by leading to a different allocation of attention across the probe array and potentially therefore also biasing encoding strategies. However, evidence from experiments 2 and 3 suggests that, though such effects might exist, they are not sufficient to account for the apparent lack of configural difference effects in either presentation condition.

A further difference between the current design and some previous studies showing configuration effects using simultaneous presentation (Vidal et al., 2005) is that in the 'change' condition in experiments in this chapter, only the spatial locations of each item change, but the values within the overall array remain constant (unless the target item changes, then it has a new value that was not in the original array). This means that when the target itself does not change, statistical summaries such as the ‘average’ or ‘variability’ would be identical between memory and test arrays. It is unclear exactly which mechanisms (such as ensemble encoding) underly the configuration effect but if such effects do contribute in some way then future research may benefit from making the ‘change’ condition more distinct with the non-cued items not only shuffling their respective locations but also changing identities. This more distinct condition might help to maximise the detection of configuration effects, especially in conditions where they may be weaker, as they may be with sequential presentation, and if such effects still fail to be observed then it may provide more convincing evidence of their absence. However, previous studies have also shown configuration effects even when using designs similar to those in this chapter, where the non-cued items merely change positions rather than introduce new values (Jiang et al., 2000), so this difference cannot account entirely for the null results.

Overall it is still unclear what may have caused the null results even in the simultaneous presentation conditions. Following experiment 1, the two likely candidates seemed the use of the retro-cue and the complex mixed block designs. However, experiments 2 and 3 found no evidence to support these explanations. The consistent lack of differences based on configuration across all conditions and all three experiments makes it unlikely to be due to a spurious sample. The sample size was above, or similar to, sample sizes used in previous experiments. It might simply have been a specific combination of stimulus parameters and, if so, this might indicate the effect does not generalise as robustly as previously thought. However, this would be surprising, as the configuration effect itself seems to have been observed across a range of different parameters and research groups. Future research aiming to establish whether such effects can occur across time might benefit from attempting to replicate the exact parameters used in previous research to eliminate this explanation, but wider research may also benefit from exploring which factors affect or limit the mechanisms underlying configuration-like effects.

As already stated, the consistent lack of any configuration effects in the simultaneous presentation conditions in this chapter make interpretation of the lack of effects in the sequential presentation conditions more ambiguous. Future research should continue to seek to replicate configural effects under simultaneous presentation conditions and then attempt to make a direct comparison with otherwise identical parameters but sequential presentation. The presence or absence of such effects in the sequential condition would then allow for much more robust conclusions to be drawn. However, as outlined above, the presence of such effects in the simultaneous condition should not be considered an absolute prerequisite. It is feasible that conditions might exist under which configural effects can be observed in a sequential but not simultaneous presentation method, and effects due to changes in temporal configuration (the temporal position or temporal order that stimuli were displayed, regardless of their spatial position) would not be expected to produce similar effects in the simultaneous condition anyway. Therefore, with the mentioned caveats, it remains necessary still to discuss potential other reasons for the lack of effects in the sequential presentation condition, and to make comparisons with previous studies that did show configuration effects, though further research would be required to investigate these possibilities further.

One possible explanation for this lack of configuration effects for sequentially presented items might be the relative differences between the spatial and temporal resolutions that the visual system delivers. As outlined in the general introduction, there are many levels of representation within the visual system and these can range from very coarse early summaries of the overall scene to later more precise and detailed representations of individual objects at more advanced stages. Much of this processing, including perhaps processes such as ensemble encoding, are likely to pool and combine inputs that are present simultaneously, and perhaps similar mechanisms do not operate as robustly (if at all) across temporally distinct inputs. This seemingly contradicts some of the earlier examples claiming to demonstrate contextual or configuration effects occurring across temporally distinct events (Blalock & Clegg, 2010; Olivers & Schreij, 2014; Silvis & Shapiro, 2014). However, as noted previously, there were issues with those designs that made it difficult to conclude confidently whether the behavioural differences between conditions truly were due to configuration effects.

One possible reason for this discrepancy could be that in those previous examples (Olivers & Schreij, 2014), memory performance for a given stimulus might have been affected by temporally distinct states of the same stimulus, not of other stimuli (it might not even matter if the preceding or following stimuli truly were the same stimulus, so long as they were perceived to be). This distinction would indicate that when a representation of an item is formed and subsequently updated over time, the updating process may not be absolute and the previous states of the representation may influence the new form to some extent. A similar mechanism such that the memory for an object’s status at a critical time might also be affected by the updating of perceptual representations of that same object as it continues to change over time. However, the current results perhaps suggest that representations of different objects might be kept distinct, such that the temporal updating of a representation of an object influences other representations over time of only itself and not separate objects. Similarly, other results may have reflected traditional spatial configuration effects being sensitive to the addition or subtraction of items between encoding and test arrays, and not between-item calculation occurring across time (Silvis & Shapiro, 2014).

However, studies have shown that more general ensemble encoding such as the average value or relative differences between items (not specifically configuration effects) do appear possible for sequentially presented stimuli (Albrecht & Scholl, 2010; Corbett & Oriet, 2011; Gorea et al., 2014; Haberman et al., 2009; Hubert-wallander & Boynton, 2015b; Jiang & Kumar, 2004; Kool et al., 2014; Ryan & Villate, 2009). As has been noted previously, the exact mechanisms underlying the configuration effect are not fully understood and the lack of configuration effects in the current studies may suggest that the ability for relative or ensemble encoding are not sufficient for configuration effects, and though the former may be possible across temporally distinct stimuli, the latter may not be. However, many of those studies employed designs where participants were required to make explicit judgements or responses concerning the summary statistics themselves, perhaps leading to artificial or purposeful strategies. Part of the aim of the current design was for the task to require the detection of only a single item, such that the similarity of the non-cued items should have had a low task relevance and any influence be more implicit and automatic. The lack of any such configuration effects in the current results may therefore suggest that they do not occur naturally across sequential stimuli and previous evidence for the ability for such encoding are dependent on specific designs where such calculations are explicitly prioritised.

Instead, much relative between-item information may not be stored as distinct representations (the capacity of such explicit representations is seemingly too limited for the vast amount of relative information to be stored explicitly) and might instead be epiphenomenal side-products of other representations rather than explicit separate representations. This interpretation might also predict that such effects would be observed more reliably across space than across time. As noted previously, there are many levels of representation within the visual system and representations generally proceed from coarse broad levels in early stages to precise detailed information at later stages. Importantly, even when those later stages are reached, the maintenance of these advanced information may depend also on the continued maintenance of some of the earlier stages as a foundation basis. It might be this foundation maintenance that underlies most of the group-level effects observed in previous studies. Importantly, there is not an obvious equivalent progression of temporal processing of stimuli in the visual system, necessarily progressing from low temporal resolution to high temporal resolution. This lack of earlier ‘coarse’ temporal representations supporting higher resolution information may therefore explain why group level representations appear to be less readily formed over time than over space.

It is possible that such temporal effects may exist but may depend on different parameters than used in the current research. Most likely is the possibility that any such ‘coarse’ representations may operate over still much shorter timescales than used in the current studies. Future research could investigate whether temporal group-level memory effects can be observed if stimuli are presented within shorter time periods, but this will raise the challenge both of task difficulty and risking that any effects are more perceptual in nature (whereby the different stimuli may be combined and perceived as one stimulus, so any supposed ‘between-objects’ effects may reflect these perceptual combination processes rather than the formation of representations of relative information between distinct objects).

One approach that future studies might benefit from adopting would be to remove the spatial aspect entirely and instead use displays of sequentially presented stimuli, each at the same single location. This would move further from the previously established designs in the spatial domain but would remove one of the potential issues of the current presentation method, where the unpredictable location of each item in the sequential condition may have disrupted efficient attentional deployment, which in turn may have weakened potential group level effects. Such a design was considered when planning Experiment 1 but the potential difficulty seemed to be in how to cue reliably a specific target if still using a partial report paradigm. However, the partial report measure may not always be necessary. Other studies investigating the configuration effect have sometimes required participants not detect the change of one specific item, but if the probe array contains any new values relative to the original memory array, with the comparison primarily between when array items maintain their original locations or change locations (e.g. Treisman & Zhang, 2006). Adopting such a design, where the task requires detecting any change, rather than one specific change, may also allow for the inclusion of a greater number of stimuli in the array, which would be expected to strengthen the utility of any group level summary statistics (Robitaille & Harris, 2011). For example, the task may require participants simply to detect whether a new item is presented in the second array, with the other items either presented in the same temporal order or a different temporal order, or participants could be shown arrays of different items but asked to detect changes in summary statistics such as a difference in the average value between the two arrays. Either of such effect would provide evidence of such group level effects operating across sequentially presented stimuli.

It should also be noted that some parameter choices in the current studies might have been expected to increase the probabilities of observing such between-item temporal effects. For example, there was no blank interval between successive items in the sequentially presented stimuli. This means that at the moment one item is removed from the screen the next item was presented. Although technically therefore not presented at the same time, lags such as phosphor decay of the monitor itself or perceptual lags such as iconic memory, may have meant that for a very short period, more than one of the stimuli may have had at least some level of visibility simultaneously. Even with this, no such effects were observed. It might therefore also be noted that such low-level transient visual information may not be sufficient for between-item associations to form, and these may instead require sustained and focussed processing of stable visual scenes.

One further weakness of the current design was that the exemplars used in the memory task were easily nameable colour categories such as “red” or “blue”. Participants might therefore have adopted a strategy where they attempt to memorise the array not visually but as a verbal list of colour names. Such a representation would lack the critical low level visual features such as ensemble encoding and so have weakened any configuration dependent effects.

Configuration effects have been observed previously even when using similar stimuli types (Jiang et al., 2000), but this does not negate the potential for such strategies here. Common methods to counter this possibility are to use stimuli that are less easily labelled verbally (such as different shades of a single colour category) or to employ an articulatory suppression task (as has been used in other previous studies reporting positive effects; Vidal et al., 2005). However, if adopting the former, future research may need to investigate how configuration effects are affected by such an increase in homogeneity among the stimuli.

A final issue that needs addressing is whether the null results reported here could be due to a lack of sufficient power. Sample sizes were designed to be either similar to or greater than many previous studies with similar designs and reporting the expected effects. However, even if positive effects have been reported with such sample sizes, they may still lack power to do so reliably, potentially leading to false negative results in subsequent studies.

Previous comparable studies often demonstrate consistently high effect sizes (often η_p^2 of at least .6; e.g. Jiang et al., 2000; Mutluturk & Boduroglu, 2014; Vidal et al., 2005) when using static simultaneous displays. An analysis using G*Power (Erdfelder, Faul, Buchner, & Lang, 2009) confirms that to detect effects of this size reliably (with a power of 0.8) would have required only 9 participants (perhaps accounting for the often low numbers of participants in previous studies). Therefore, the failure to observe configuration effects even when using only static displays (as in experiment 2) cannot be accounted for simply as a lack of sufficient power. The equivalent effect sizes reported here instead were much smaller (η_p^2=.138 to .143), suggesting that the null results were likely due to a genuine absence or weakening of the predicted effects, perhaps due to the discussed methodological features.

For the sequential array conditions, it is more difficult to make direct comparisons with previous studies, but focussing on those that report some configuration-like effects while using dynamic or sequential stimuli show a wider range of effect sizes, including η_p^2=.243 (Olivers & Schreij, 2014), η_p^2=.4 (Silvis & Shapiro, 2014) and η_p^2=.65 (Papenmeier et al., 2012). Sample sizes of between 14 (for an effect size of .4) and 27 (for an effect size of .25) would have been needed to detect the lower end of this range. The actual effect size observed in the sequential only condition of experiment 3 was a little weaker than even the lower limit of the range reported (approximately η_p^2=.2), so it might be argued that the current sample sizes of 20 would have lacked sufficient power to detect such effects. However, even this marginal effect was driven by differences between the isolation condition and both full array conditions, with the full array conditions showing almost identical performance regardless of whether the configuration had changed or not. It therefore remains unlikely that the lack significant differences between the configuration conditions specifically could be accounted for simply by a lack of power

Having failed to replicate even previously established spatial configuration effects using our chosen parameters, it was decided not to pursue the more advanced and more sensitive initial questions, such as the consequence of changing the temporal order but not spatial position of items, or whether the expected configuration effects are more sensitive to changes temporally closer to the critical time window than those much earlier/later. Instead, the next chapter will investigate instead how the spatial configuration effects operate within the tactile modality. This was chosen as the tactile modality is perhaps even more likely to exhibit across region integration than vision, and so such group level representations may play an even more critical role.

Chapter 3: Configural Effects in the Tactile Domain

General Introduction

So far, hierarchical representations have been discussed only in relation to vision. However, it is highly unlikely that these effects are unique to vision only. In everyday life, especially when using our fingers to feel the texture of a surface, we tend to integrate the inputs from multiple fingers to form an overall texture (and likely discount fingers that are suspended in the air, not in contact with the surface), rather than feel different surfaces simultaneously with different fingers. Therefore, any effects that depend on the integration of inputs across fingers, such as calculation of a 'mean' or relative encoding, might be expected to be even greater in the tactile modality.

This chapter describes two experiments that investigate whether the configuration effect previously described in the visual domain (Jiang et al., 2000; Vidal et al., 2005) may generalise to the tactile domain, and, if so, whether these effects show any qualitative differences between the two sensory modalities. Furthermore, as will be discussed below, most research into tactile memory has focussed on the use of vibrotactile stimulation (delivered by small vibrating devices that can be attached to various locations of the body), with the critical feature to be remembered being the location on the body that was stimulated. However, there is evidence that this form of tactile stimulation may behave qualitatively differently to texture perception, so an additional focus of this research will be to investigate memory performance for tactile textures.

Existing studies of tactile working memory suggest that participants with visual working memory are more likely to exhibit higher tactile working memory, but that tactile working memory is poorer and more variable than its visual counterpart (Bliss & Hämäläinen, 2005). We appear less able to 'subitize' (rapidly and accurately count the number of presented stimuli) using touch (up to roughly three items; Gallace, Tan, Haggard, & Spence, 2008) than vision (up to four items; Trick & Pylyshyn, 1994). This might indicate that touch has a lower threshold for separate stimuli to be 'grouped' rather than represented as independent and distinct stimuli, perhaps suggesting that configuration effects will be particularly strong in the tactile domain. Other than just subitizing, when using vibrotactile stimuli, participants are only able to recall correctly the locations of up to approximately three stimuli after long delays (one second or greater; Auvray, Gallace, & Spence, 2011), with change detection sensitivity dropping even further if transient tactile masks (Gallace, Tan, & Spence, 2006) or even visual masks (Gallace, Auvray, Tan, & Spence, 2006) are introduced. Similar drops in change detection performance can be observed even if no such transient is presented but instead a bodily movement is made (Gallace, Zeeden, Röder, & Spence, 2010).

Together these results suggest that upper estimates of tactile memory may reflect fragile memory representations, and the number of items that can be stored as true 'robust' working memory-like representations that can survive secondary tasks or intervening masks, is even more limited than in vision. Such a low performance is surprising given the task is often simply to detect and localise stimulation (rather than identify or discriminate different forms of stimulation). Some of these results may be explained by short 200ms stimulus durations (Auvray et al., 2011; Gallace, Tan, et al., 2006), such that representations may suffer from masking or noise from incomplete encoding. However, when an intervening mask is used, performance drops markedly even when three stimuli were presented for 1000ms (Gallace, Auvray, et al., 2006). This may suggest that, unlike in vision, even just three stimuli are not encoded into a robust working memory representation within 1000ms.

Interestingly, even though participants are poor at detecting specific changes in stimulation, if required to give a forced choice between whether any change that trial was for increased or decreased number of stimuli, participants did perform above chance, suggesting they may have some hierarchical representation of the change but not sufficient to reach an explicit 'change' response threshold (Pritchett, Gallace, & Spence, 2011). Furthermore, this relatively low 'capacity' for maintaining separate items in a tactile working-like memory could feasibly allow for a relatively greater role for other levels of information, such as the integration of group level statistics, in a tactile memory task.

However, these studies have so far focused on vibrotactile stimuli, with the critical feature being location of stimulation on the body. It is unclear whether such results would also generalize to the encoding and maintenance of different surface textures and, rather than simply where stimulation was, what the stimulation was. One example of how paradigms using real textures may yield different results than fixed vibrating stimuli is that some features such as the subjective perception of surface roughness seems more critically related to the spatial distribution of stimulated receptors, rather than simply changes in the timing or frequency of such activations (Connor & Johnson, 1992), and therefore a feature dimension that cannot be well simulated using fixed vibrating stimuli. One highly relevant example found that practice identifying texture stimuli has been shown to lead to subsequent improvements not only on the trained finger, but also neighbouring fingers (Harris, Harris, & Diamond, 2001). Importantly, this result did not occur when vibrotactile stimuli were used, suggesting there may be greater integration across fingers for the perception of texture than the stimuli used in most previously reported experiments.

Further evidence of such integration across fingers is that an irrelevant texture presented on one finger has been shown to affect the perception of a target texture on another finger, such that a rough distractor would make the target appear more rough, and vice versa (Kahrimanovic, Bergmann Tiest, & Kappers, 2009; Roberts & Humphreys, 2010). Similarly, following prolonged tactile stimulation, resulting ‘aftereffects’ have been shown to transfer also to other fingers (Van Der Horst et al., 2008). Such integration effects occur both between fingers on the same hand and across hands (Roberts, 2013; Van Der Horst et al., 2008). Together, such results indicate that information from disparate tactile inputs show at least some integration. This integration of information across inputs cannot be accounted for purely by physical limits of perception, such as an automatic summation of information across receptive fields, as the magnitude of this integration seems to vary between different conditions. Instead, this integration, with a biasing of individual inputs towards the mean value, may be analogous to similar effects previously observed in vision (Brady & Alvarez, 2011). Finally, such integration has been shown to occur not only spatially adjacent locations, such as fingers on the same hand, but also (to a reduced but still reliable level) across different hands (Evans, Craig, & Rinker, 1992), lending further support to the possibility that group level information may play an important role in tactile perception. This is the focus of the current chapter.

Although previous research using vibrotactile stimuli has used groups of stimuli (often up to six), research into texture has often focussed either on the perception of a single texture at a time (as is most common in everyday life), or how pairs of textures (simultaneous or sequential) may influence perception of each other. The current studies use a larger group of texture stimuli, allowing for potentially more complex or nuanced mechanisms – such as the derivation of summary statistics or the influence of relative (rather than absolute) encoding – to arise. It is also hoped that using larger groups of textures will allow us to measure working-memory-like performance for tactile texture stimuli.

Experiment 4: Configuration Effects in the Tactile Domain

Introduction

The initial experiment aims simply to examine whether the configuration effects previously reported in the visual domain (Jiang et al., 2000; Vidal et al., 2005), as outlined in Chapter 3, can also be observed within the tactile domain. Participants will encode multiple tactile stimuli but will only be asked to detect the presence or absence of change of a single input on any given trial (retro-cued at the end of the memory maintenance period, still requiring participants to encode and maintain all presented stimuli up until that point). The key prediction is that if group level representations, such as ensemble encoding or relative values between items, play important roles also in tactile texture perception, then memory for the cued target should be improved on trials where the non-cued inputs are also presented during the test phase with identical properties as in the encoding phase, relative to when a different combination of stimuli are presented to the non-cued locations. This is because if the non-target items retain their same values and configuration as in the encoding phase then participants may be able to use any detected changes in the summary statistics between the two arrays to infer that the target must have changed, even if they do not reliably remember the specific original value of the cued location.

Method

Participants

Twenty participants (2 male, 18 female) aged between 18 and 38 (M=21.3, SD=5.1) took part in the study in return either for course credit or paid compensation. This number was comparable or greater than previous studies investigating working memory for tactile stimuli (e.g. 11 participants; Gallace, Tan, Haggard, et al., 2008; Gallace, Tan, & Spence, 2008). 15 were right-handed. All procedures were reviewed and approved by the Departmental Ethics Committee.

Stimuli and Materials

Each stimulus trio consisted of three ribbons glued to a rectangle of card, with a perpendicular cardboard ‘handle’ on the back, allowing the experimenter to grip and move each stimulus from above. There were five differently textured ribbons and 27 cardboard panels. Each cardboard panel held three different textures (no panel held two ribbons of the same texture) and could be rotated 180° to present the three ribbons a reversed order. This allowed for a total of 54 total possible combinations of three ribbons (trios), though note that some panels contained the same three textures in the same order, leading to some duplicates. There were therefore 40 unique combinations of textures possible.

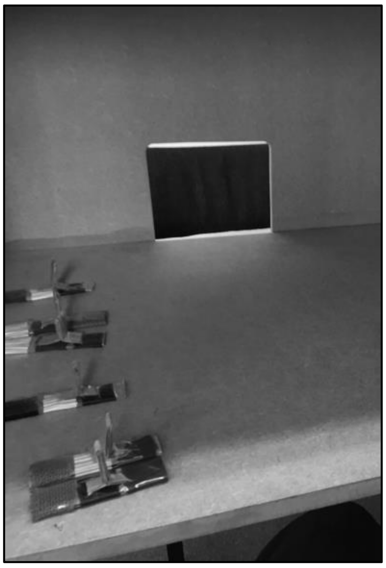

Figure 2: Example of view from experimenter's side, showing one of the holes in box that participants would insert hands through, and examples of some of the stimuli used, each with three textured ribbons glued to card panels.

An algorithm was used to select the stimulus trios used in each trial. First it would select two stimulus trios for the encoding phase at random. Then, for the two trios to be used in the probe phase, all possible combinations were selected where the target finger for that trial was either the same or different (depending on the condition for that trial) and where every non cued position was either the same or different (depending on the condition for that trial). If no such stimuli were possible then the algorithm would restart and select a different two random stimuli to use during the encoding period, and then check possibilities for the probe phase again.

Of the possible trios that could be used for the probe phase, two were selected at random. Note that the only two conditions were that the texture on the target finger must have been either the same or different (depending on the current condition) to the memory phase, and that every non-cued item should have been either the same as or different (depending on condition) from those that were used in the memory phase. No conditions were enforced regarding the number of repetitions or how they should change (if they changed). Importantly, even when all textures were to stay the same, the actual trio used could differ if there was more than one trio that contained that specific combination of textures. This meant that even when textures were to remain the same, two physically different but otherwise duplicated trios might have been used.

The experiment was carried out inside a large wooden ‘box’ that was positioned between the participant and the experimenter. This box had two holes cut out of the bottom of one panel (that facing the participant), for participants to put their arms through, while the opposite panel (facing the experimenter) was missing such that the experimenter could fully access the interior of the box. Throughout the experiment participants would insert both arms through the holes and position their hands such that palms were facing upwards. Sponge cushions were provided for them to rest their forearms on, and the fingers of the hands would rest upon a pouch filled with rice. This enabled participants to press down such that the rice would mould around their fingers, helping to keep them secure and in an approximately level horizontal plane. One issue if this were not present is that participants had a tendency over time for their fingers to curl inwards and often one finger would end up ‘lower’ than the other two, meaning that when the texture panels were presented, the lower finger would not be sufficiently stimulated.

The target was cued by a gentle tap on the very tip of the relevant finger by a cotton-wool bud. Throughout the experiment the experimenter also wore latex gloves, and both experimenter and participant used anti-bacterial hand gel at the start of each session.

Procedure

On each trial, the algorithm selected and displayed the identifiers of the stimuli required for that trial to the experimenter. The experimenter then got all four of these ready at the start of each trial. On any trials where only two or three stimuli were required (such as when the same stimuli were used in both the encoding and probe phases), the experimenter still selected additional random stimuli and placed them perpendicular to the true to-be-used stimuli (to avoid the experimenter mistaking them for valid to be used stimuli), in order to prevent any auditory clues for the participant, from counting the number of items selected.

Each trial contained four phases, an encoding phase (where the stimuli to be remembered were presented), a maintenance phase (where nothing happened but waiting for a specified amount of time to elapse), a probe phase (where the probe stimuli were presented) and a response phase (where the experimenter waited for the participant to make their response). Figure 3 illustrates an example trial for one of the hands (though note that stimuli presentation was always on both hands simultaneously).

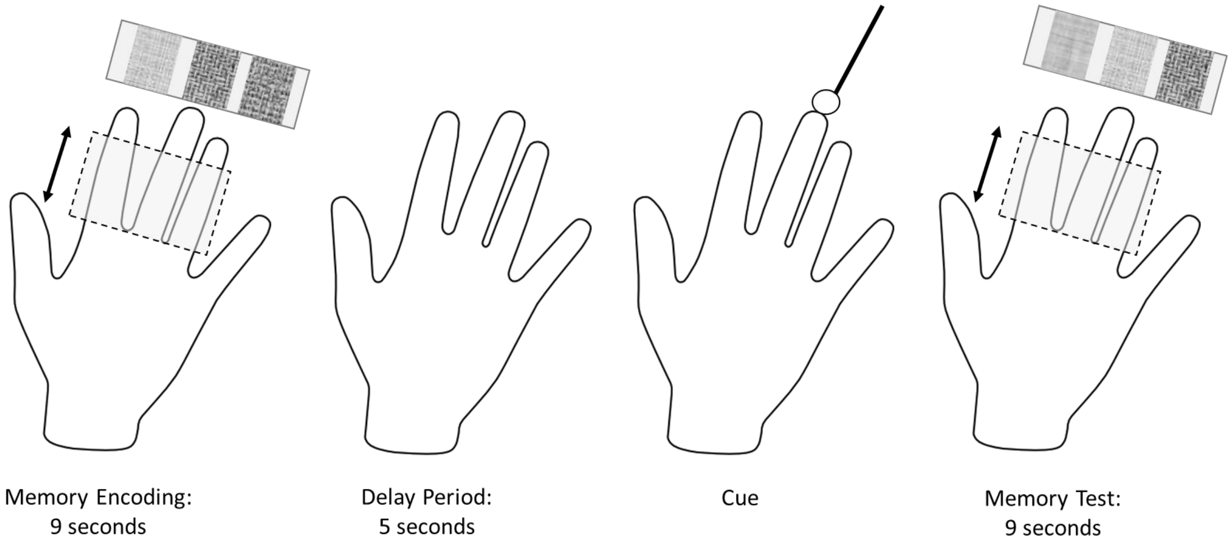

Figure 3: Schematic illustrating example trial progression. Figure shows only left hand but stimulus presentation was always to both hands simultaneously. In example trial, middle finger is cued as target, the cued texture changes and also the non-cued textures change.

Participants sat opposite the experimenter but each was unable to see the other due to the large wooden box between them. Participants would insert their hands through the holes in the box and rest their forearms on the sponge, and their fingers (palms facing upwards) on the rice-filled pouches. The experimenter checked that the fingers were spread sufficiently far apart such that no individual ribbon should stimulate two fingers simultaneously, and would help to ensure fingers were stable and on a similar plane, such that the texture panels would be able to stimulate all three middle fingers simultaneously. Participants were shown this setup before the start of the experiment to reassure them and were given the opportunity between every trial to rest their arms and hands if they wished.

When both the experimenter and participant were ready, the experimenter initiated the trial by pressing a foot pedal. This caused a visual cue to appear on the experimenter’s screen (not visible to the participant), which was programmed to move at the desired rate of movement of the tactile stimuli. The experimenter then stroked the two stimuli selected for the encoding phase along the participant’s middle three fingers, trying to keep at roughly the speed of the visual cue, while avoiding stimulating the fingertips themselves (as these were thought likely to be too sensitive which could lead to ceiling effects). Note that the medial texture (closest to the torso) was presented to the ring finger of each hand, the middle texture was presented to the middle finger of each hand, and the lateral texture (closest to the thumbs) was presented to the index finger of each hand). This stimulation lasted for nine seconds, reversing direction every three seconds.

Following the end of the encoding stimulation, the visual cue vanished, and a timer counted down for five seconds. This was the maintenance period, and no stimulation was applied (both experimenter and participant simply sat motionless). Following the end of the five second maintenance period, the experimenter tapped the fingertip of the target finger with the cotton-wool bud and, as quickly as possible, began to stroke the stimuli selected for the probe phase across the participants fingers. Again, this stimulation was accompanied by a visual cue for the experimenter and persisted for nine seconds (changing direction every three seconds). Note that there was some variability in the transition from the maintenance period to the probe period as, once the countdown finished, the experimenter would need to cue the target finger and then quickly pick up and position the probe stimuli. This jitter means that the true duration of the ISI between the offset of the encoding stimuli and onset of the probe stimuli ranged generally between five and six seconds.

Following the probe phase, stimuli were removed from the fingers and the experimenter would wait for the participant to give their response verbally ("same" or "different"), before then pushing one of two foot pedals to log which response had been made. Participants were instructed to prioritise accuracy and that there was no advantage to responding quickly. Once a response was recorded, the experimenter asked the participant to indicate when they were ready for the next trial. If the participant gave their response during the probe phase (rather than waiting until after) then the full stimulation was still given, allowing them the same total time to consider their response as other participants, including the option to change their response if desired. There was no time limit for responses, though all were given within a few seconds.

Every participant first completed three practice trials, where the experimenter talked through the procedure and answer any questions or clarifications the participant had. Following this short practice block, each participant then completed two blocks of twelve trials each. Within each trial were six trials where the non-cued items were the same between the encoding and probe phases, and six trials where they changed (giving a total of twelve for each condition for each participant).

Note that the number of trials on which finger was the cued finger, whether the target changed or not, and whether the non-cued targets changed or not, were fully counterbalanced, such that there were an equal number of trials for each combination of conditions.

Design

A repeated measures design was used with a single factor of whether the non-target items in the probe phase were the same or different as the non-targets at encoding. The dependent variables were accuracy (% correct), memory capacity (Cowan's k), d' (signal detection theory) and criterion (signal detection theory) of correctly identifying whether the target texture had changed or not between the encoding and probe phases.

Data Analyses

Given the manual nature of the paradigm, there were occasional trials in which the experimenter made an error. Unfortunately, not enough information was recorded to allow these to be safely recoded manually and they were instead excluded. In total 0.4% of trials were excluded in this way and all subsequent percentages are reported in relation to the number of included trials for that condition (no participant had more than one trial excluded).

Following any exclusions, accuracy was calculated for each condition as the number of correct responses divided by the maximum number of correct responses for that condition for that participant. Any participants whose average performance across both conditions was 50% or less was then excluded (two participants were excluded in this way, leaving a final sample size of 18).

Memory capacity was calculated using Cowan's k (Cowan et al., 2005; note that although the probe array did include 6 items, only a single one of those items was actually being probed; Rouder, Morey, Morey, & Cowan, 2011), and was calculated as $k = N(H – FA)$, where H is the hit rate, FA is the false alarm rate, and N is the set size of the display.

D' scores were calculated as $d' = z(H) – z(FA)$, and criterion was calculated as $c = -0.5 \times (z(H) + z(FA))$, where H is the hit-rate, FA is the false alarm rate, and z is the normalised score. Note that corrections were applied where $H=1$, transforming it to $H=1-(1/(2*T))$, and where $FA=0$, transforming it to $FA=1/(2*T)$, where T is the number of trials (12).

Separate paired sample t-tests were then run on the accuracy and d' scores, to test for any differences between when the non-cued items were the same versus when they were different across the encoding and probe stages of presentation.

Results

Overall, performance seemed slightly higher when the non-cued items were the same in the probe phase (M=70.0%, SD=13.7) than when they were the different (M=64.8%, SD=15.5), but this difference was not statistically reliable, *t*(17)=1.084, *p*=.294, d_z=0.256. Figure 4 illustrates the overall pattern of results.

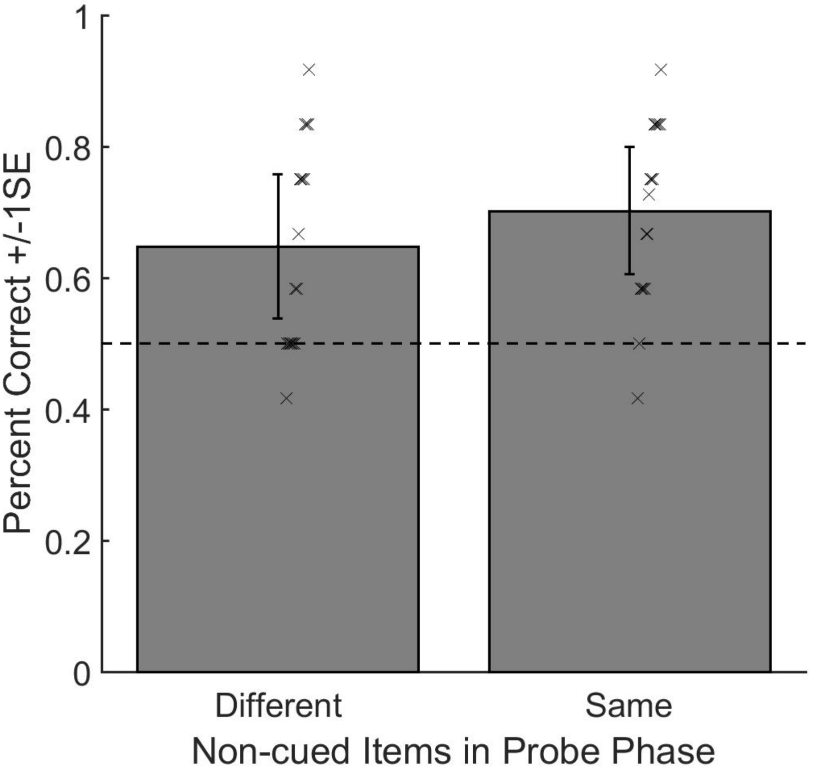

Figure 4: Graph showing accuracy (% correct) of cued target change detection, depending on whether non target textures had same or different values/configuration. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

Estimated memory capacity also seemed to be slightly higher when the non-cued items were the same (M=2.34, SD=1.57) than when they were different (M=1.75, SD=1.82), but this difference was also found not to be statistically reliable, *t*(17)=1.84, *p*=.334, d_z=0.433.

Similarly, although sensitivity (d') appeared slightly greater when the non-cued items were the same (M=1.24, SD=0.88) than different (M=0.87, SD=0.95), this difference was not statistically reliable, *t*(17)=1.26, *p*=.224, d_z=0.297.

However, criterion (bias towards responding in one way or another) was found to be statistically different between the two conditions, *t*(17)=3.01, *p*=.008, d_z=0.709. Participants were significantly more likely to state that the target had changed when the non-cued items had not changed (M=-0.30, SD=0.37) than when the non-cued items had changed (M=0.07, SD=0.36).

Discussion

Overall, the current study found no reliable evidence of any difference in memory performance as a function of whether the non-cued items in the probe phase were the same or different to those in the encoding phase. It is worth noting that the direction of results was in the direction predicted (better performance when the items were the same) but the current data do not provide sufficiently reliable evidence for any such conclusions to be drawn.

When looking at the raw data, many participants were observed to be at (or below) chance in the ‘different’ condition (and some also in the ‘same’ condition). Although these participants performed above chance in the other condition, giving them an overall mean performance above chance, it appears that many participants struggled to perform the task accurately. Indeed, only nine participants had scores above 50% in both conditions.

Interestingly, in line with these observations, the measures of memory capacity in the current study were noticeably lower than previous estimates. Studies using vibrotactile stimuli have estimated a capacity for just three stimuli reliably (though perhaps up to five if using a partial report paradigm with a sufficiently short memory delay interval, Gallace et al., 2008) while the current results suggested a capacity of around two stimuli on the texture memory task used here. It is not clear whether this difference between the current results and earlier findings is reliable or not, and future research may wish to compare capacity more directly between different types of tactile stimuli. However, if such a difference is confirmed then the smaller capacity for texture stimuli may be surprising given that previous results in the visual domain have suggested that more realistic stimuli often yield greater memory performance than abstract stimuli (Brady et al., 2016). The advantage of real-word objects that has been seen in the visual domain is thought to be partly due to these objects providing richer structured information at various hierarchical levels of representation than more homogenous unidimensional stimuli (Brady et al., 2016). This would have been another reason to expect that capacity would have been higher for the textures than for simple individual vibrations, because there are multiple dimensions along which changes in texture might be detected (e.g. rough/smooth, but also firm/soft). However, this expectation was not be supported by the lower capacity that was observed here for texture stimuli. Although the real-life textile textures presented to the fingers may be argued to be more ‘realistic’ than vibrating tactors at various body locations, the latter might be argued to be more similar to stimuli such as a fly landing on one’s skin or a phone vibrating in a pocket. One might also question how often a set of six differing textures would need to be processed and encoded simultaneously in real world conditions.

One alternative explanation for this difference could instead concern the nature of the memory task (rather than the nature of the stimuli themselves). In the previous studies mentioned (e.g. Auvray, Gallace, Hartcher-O'Brien, Tan, & Spence, 2008; Gallace, Auvray, et al., 2006; Gallace, Tan, Haggard, et al., 2008; Gallace, Tan, & Spence, 2008), the critical feature to be remembered has typically been the location of otherwise simple and often identical stimuli, so participants had to remember 'where' stimuli were. In contrast, participants in the current experiment were asked to remember 'what' stimuli were. Given the hierarchical nature of perception and memory (Brady & Alvarez, 2011; Juan & Walsh, 2003; Kimchi, 2015), it should not be surprising that difference in performance may be observed at different levels of representation. In other domains such as vision, there are many examples of results where participants might be able to report that a stimulus was present (where) without necessarily being able to report accurately its identity (what; Haque, Lobier, Palva, & Palva, 2020).

To understand better which differences in the stimuli may affect most differences in memory performance, future research may consider a more direct comparison using both types of stimuli. For example, a location-based memory task such as the previous studies outlined may be performed but using texture stimuli rather than vibrotactile stimuli (though the practicalities of such a design may prove difficult). Alternatively, one could repeat the current paradigm but using vibrotactile stimuli that differ along one or more dimensions (such as vibration frequency or intensity).

An interesting and unexpected result was the apparent difference in criterion between the two conditions. Participants were more likely to report that the target had changed when the non-cued items did not change compared to when they also changed. This is the opposite direction than might have been expected. For example, on a trial where the participant has no memory at all of the cued target and so is going to knowingly just guess, it might have been expected that they may have based their guess on whether they detected a change or not in the non-cued items. Such a strategy would not have offered any performance advantage (such a guess would have been equally likely to be correct or incorrect), but it might have been an unsurprising strategy when they had nothing else to base their guess on.

This difference seems to be driven more by an increased bias when the non-cued items remained the same rather than when they changed (the criterion shift is closer to zero in the latter condition than the former). One possible explanation could be the fidelity of the memory for representations of the group at lower hierarchical levels. For example, one of the proposed mechanisms underlying the configuration effect in vision is maintaining a memory for the average (e.g. the average size). Such a representation would mean that even if a participant forgot the individual identities of each item making up the array, they might still be able to detect changes if the average value of an array changed. When the non-cued items remained the same, any such change to the average value could only occur if the cued target had changed, and so decisions based on the average value would theoretically yield results as good as though the cued item itself had been remembered. By contrast, if the non-cued items change then the average value will almost always differ, meaning that decisions based on the average alone would be at chance in that condition. Even if such trials are relatively rare, this is an example of how less precise hierarchical representations may confer some advantage to performance when the non-cued items stay the same than when they change. Generally, participants will likely usually have some memory of the item, or at least a memory of some of the non-cued items (rather than basing a memory entirely on the mean value alone). In the latter case, if participants do remember one or more of the non-cued items, then this would allow them to detect whether the non-cued items had changed or not by detecting whether even a single non-cued item changed (because either all of the non-cued items changed or none of them did). If the participant was aware that the non-cued items had changed on a particular trial, then it is possible that they would disregard any change in the average value as this would say nothing about the target item itself, whereas if they knew the non-cued items had remained the same then they may give greater weighting to a change in the average value (as such a change could only be the result of a change in the cued item). In these circumstances, the observed pattern of results could perhaps be explained by assuming that participants are not able to derive very accurate average values. In that case, on trials where the participant detects that a non-target is the same (and so all non-targets are the same) and therefore decides to use the average value in determining their response, their inaccurately maintained average may be unlikely to match the average value in the probe phase, making them more likely to report that the target has changed. This bias towards ‘change’ responses on trials where the non-targets remained the same would not necessarily affect accuracy performance, as such erroneous ‘change’ responses would be equally likely to be correct as incorrect.

Such an explanation would require that ensemble encoding mechanisms are particularly poor under these conditions which, though would be surprising given the rationale for the study, might explain the overall lack of any reliable configuration-based effects. Much of the research outlined in the introduction might suggest that integration across fingers may lead to better representation at more coarse levels of representation. However, much of the evidence, especially for textured stimuli, focussed on texture perception. It is possible that though such integration may form an important role in immediate perception but that such information cannot be maintained reliably once the stimulation is removed.

Indeed, it is possible that in the current paradigm even the observed performance may at least partially be due to verbal encoding rather than truly a memory of the textures themselves. The different textures do not form obvious distinct categories in the same way that colours might, but participants may still give them verbal labels such as the “really rough one”. This possibility is aided by the fact that there were only five exemplars in the whole session, so even if such vague labels were used, not many would be needed (such a labelling strategy would seem less feasible if twenty different texture exemplars were used, for example). Future research may consider reducing the possible confound of articulatory rehearsal, such as by changing the nature of the stimuli (such as increasing the number of different exemplars) or requiring participants to perform a concurrent articulatory suppression task.

Although in general measures of working memory capacity could be considered relatively independent of overall set size, the lower-than-expected capacity may also account for the overall poor performance. Given the current set-size of six items, if capacity had been ~3 items (as previous results suggested) then an average performance of approximately 75% would have been expected, but a capacity of two would be expected to yield a performance of ~67% (similar to the means observed in the current experiment). However, these values of means and individuals may differ. For example, if a participant had a memory capacity of 1 item, the chances of that remembered item being the probed target would be just 1/6, or two trials per condition. Considering that each trial equates to ~8% (so even small variability may lead to seemingly large percentage variability), and considering random noise or variability, it may have been difficult to distinguish participants who were genuinely guessing (or simply unable to do the task) from those who performed badly but may have averaged above chance over a greater number of trials. However, even beyond mathematical issues with these low trial numbers, the apparent difficulty of the task may also have led to goal neglect (Duncan et al., 2008), where participants may have disengaged. Especially given that feedback was provided after every response, if too many responses were incorrect then they may have decided the task was too hard and stopped trying.

Given the above considerations, Experiment 5 will reduce the set size to only four items presented simultaneously (though still the same five possible textures overall). This should still exceed the apparent memory capacity observed here while nevertheless reducing the difficulty of the task. This would improve overall performance (making it easier to distinguish between participants who performed badly from those who genuinely were at chance) and would be overtly easier, perhaps improving engagement and motivation.

Experiment 5: Configuration Effect in the Tactile Domain with Easier Task

Introduction

The qualitative pattern of results of the previous experiment were as predicted, but they lacked statistical reliability. One possible factor may have been the overall poor performance of participants and so the current experiment repeated the same paradigm but made the task simpler by using only four exemplars per trial instead of the previous six. This is achieved by presenting textures to only two fingers on each hand. Overall predictions remain the same as in Experiment 4, that memory performance for the cued target will be better when the non-cued textures remain consistent between the encoding and test phases than when they are changed.

Method

Participants

Twenty-two participants (7 male, 15 female) aged between 18 and 51 (M=23.5, SD=7.1) took part in the study in return either for course credit or paid compensation. 20 were right-handed. All procedures were reviewed and approved by the Departmental Ethics Committee.

Stimuli and Materials

Everything was identical to the previous experiment with one exception. Instead of presenting six textures per trial, only four were presented. However, the same trios were used, which each still had three textures attached. Therefore, the lower number was achieved by moving each trio laterally by one position. Whereas in the previous version, the trio was presented such that the medial texture (texture closest to the torso) was presented to the ring finger, the middle texture to the middle finger and the lateral texture (texture closest to the thumbs) to the index finger, these were shifted such that the medial texture was presented to the middle finger, the middle texture was presented to the index finger, and the lateral texture did not touch any fingers (it was suspended in space). Therefore, only the index finger and middle finger of each hand had any texture stimulation. The algorithm for selecting each trio was adjusted so that it did not consider the lateral texture (which would not be felt by participants).

Procedure

The procedure was identical to the previous experiment, with the only difference being the positioning of the texture trios, as outlined in the materials section.

Design

The design was identical to the previous experiment, with the independent variable being whether the non-cued items were the same or different between the encoding and probe phases. The dependent variables were once again accuracy, memory capacity, sensitivity, and decision criterion.

Data Analysis

Given the manual nature of the paradigm, there were again occasional trials in which the experimenter made an error and these were excluded from analysis. In total 1.5% of trials were excluded in this way and all subsequent percentages are calculated and reported in relation to the number of included trials for that participant for that condition (no participant had more than one trial excluded).

Following any exclusions, accuracy was calculated for each condition as the number of correct responses divided by the maximum number of correct responses for that condition for that participant. As in the previous experiment, any participants whose average performance across both conditions was 50% or less would have been excluded but none violated this criterion, leaving a final sample size of 22.

All measures were calculated and analysed as described for Experiment 4.

Results

Overall, performance was very similar when the non-cued items were the same in the probe phase (M=72.4%, SD=15.6) than when they were the different (M=73.6%, SD=12.7), and no significant effect was observed, *t*(21)=0.270 *p*=.790, d_z=0.058. Figure 5 illustrates the overall pattern of results.

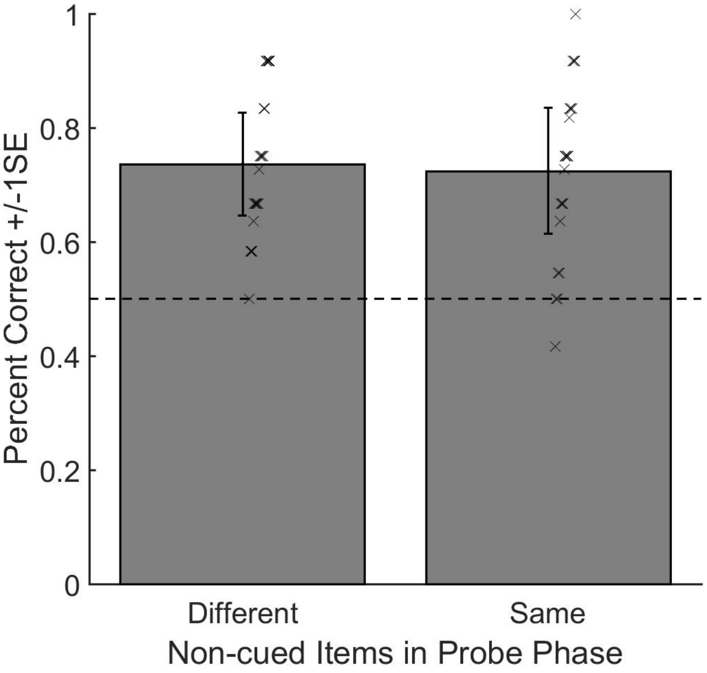

Figure 5: Graph showing performance (% correct) at detecting cued texture change, split by whether the non-cued textures had same or different values/configuration. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

Estimated memory capacity was also very similar when the non-cued items were the same (M=2.64, SD=1.77) than when they were different (M=2.70, SD=1.45), and this difference was also found not to be statistically reliable, *t*(21)=0.14, *p*=.889, d_z=0.030.

Similarly, sensitivity (d') was very similar when the non-cued items were the same (M=1.40, SD=1.01) than different (M=1.48, SD=0.86), and difference was not statistically reliable, *t*(21)=0.25, *p*=.807, d_z=0.053.

However, the criterion (bias towards responding in one way or another) analysis appeared to suggest that participants were more likely to report that the target had changed when the non-targets were the same as during the encoding phase (M=-0.231, SD=0.413) than when their configuration had changed (M=0.028, SD=0.542). This difference was found to be 'marginal' but ultimately not statistically significant, *t*(21)=2.02, *p*=.056, d_z=0.430.

Combined Analysis

To test explicitly how memory performance varied between the two experiments, the results from each were submitted to a combined 2x2 mixed ANOVA with the factors of Experiment (Experiment 4 or Experiment 5) and Configuration (Same or Different).

Overall performance was not significantly different between the two experiments (*F*(1,38)=2.988, *p*=.092, η_p^2=0.073). There was also no interaction between experiment and configuration type (*F*(1,38)=1.000, *p*=.324) and no main effect of configuration type (*F*(1,38)=0.414, *p*=.524, η_p^2=0.010).

Discussion

Overall, in line with Experiment 4, this experiment found no significant effect of whether the non-target items changed or remained the same on any of the measures of memory performance. While the data for experiment 4 matched the predicted pattern qualitatively (though were not statistically reliable), the results of the current experiment do not, and there is not even evidence of a trend towards a difference. This elimination of even the qualitative pattern of results occurred despite reducing the difficulty from 6 items to 4. This reduction in difficulty led, at best, to only a 'marginal' increase in memory performance. This may imply that the task was still too difficult, though absolute performance values seem reasonable, especially in Experiment 5, suggesting instead that the lack of effects was not simply due to task difficulty. If anything, reducing the number of targets seems to have reduced the utility of configuration effects.

Taken together, the two experiments would suggest that there is therefore no evidence for a configuration effect in tactile texture memory. It would remain possible that a different paradigm may yet yield such results, but using the current stimuli and paradigm, which were intended as approximate analogues of the visual counterparts, no such effects are seen. Initially it may be surprising that there is a lack of a configuration effect as previous studies seem to suggest a strong integration of information across the fingers, which might have been predicted to lead to even stronger configuration effects than those already observed in vision. However, as discussed in Experiment 4, much of the most relevant evidence (such as those experiments using texture stimuli) focussed on the perception of textures as they were being presented. It is possible that such information cannot be maintained effectively once the stimulation is removed. This apparent poor performance of tactile memory may not be so surprising as in daily life (or evolutionary history), it is less common for there to be a need to maintain texture information across delays. Instead, texture might more often have been most important for identifying objects while they remain present.

It is possible that some of the changes made following experiment 4, intended to increase the chances of detecting a configuration effect, may have worked against such an effect. One example might be that it is unclear just how much integration in memory would occur across the hands when a set size of four is used. Within the visual domain, it has sometimes been found that increasing (rather than reducing) set size can increase the utility and weighting given to group level hierarchical representations (Robitaille & Harris, 2011). Similarly, although tactile integration between hands has been shown, integration across fingers of the same hand appears stronger (Harris et al., 2001), and so reducing the number of stimuli per hand to two may have led to an overall weakening of such effects. Future research may consider still using only four stimuli but presenting all to the same hand rather than across hands.

It is important to emphasise that there is reason to expect that the current findings may not generalise to all areas of the skin. Especially given the textural nature of the stimuli, an important difference may exist between the fingertips and the rest of the finger (less tuned to fine discriminations or encoding of texture). Specific efforts were made by the experimenter to minimise any stimulation of the actual fingertips as it was assumed their increased sensitivity would make the task too easy. Given the overall poor performance and low memory capacity measures, this may have been overly conservative. Non-fingertips may by suitable for some types of tactile stimulation such as detecting the location of a vibration or some basic features of a texture (such as approximate roughness) but may have lacked the sensitivity to detect and encode reliably the more subtle complexities required to distinguish between the five textures used in this experiment. Future research could attempt to test this either by stimulating the fingertips themselves or, if still needing to avoid the fingertips, to investigate explicitly how distinguishable the used textures are from each other.

Some participants did mention having the sensation of still being able to ‘feel’ the textures even by the onset of the probe phase. One possibility for this, especially when comparing to results from the visual domain, may be the lack of any masking stimulus used in the current experiment. For example, in vision there is an ‘iconic’ memory that is thought to persist for the first half a second after stimulus offset (perhaps driven primarily by a gradual decline of firing of photoreceptors) but is easily eliminated by any kind of mask (Sperling, 1960). Similarly, visual ‘fragile’ working memory (Vandenbroucke et al., 2011) has a high capacity and duration, and can survive a simple light mask but not a complex noise mask (i.e. a mask with features confusable with the relevant stimuli). The five second delay period used in the current design makes any role for iconic memory unlikely (though how similar any tactile working memory may be to its visual counterpart is not well understood), though it could leave ambiguity for whether the current design was testing tactile working memory, or a form of ‘fragile’ tactile memory. Future research may benefit from testing explicitly how memory for textures is affected by the presence of intervening masks. It may be that even the relatively low memory capacity observed in these experiments may drop further under conditions that eliminate all but traditional ‘robust’ working memory.

General Discussion

Overall neither Experiment 4 nor Experiment 5 showed any evidence of a configuration effect. As mentioned in Chapters 1 and 2, the exact mechanisms underlying the configuration effect are not fully understood, even within the visual domain. However, likely candidates that are expected to play a major role are those of ensemble encoding and of relative encoding. The lack of any configuration effects in either of the experiments in this chapter might be taken to imply that such mechanisms play less of a role in tactile perception than in memory. This is surprising given existing evidence for some similar effects as those observed in vision, such as the biasing of the perception for individual inputs towards the average value (Corbett, 2017). The lack of expected differences is made even more surprising using texture stimuli rather than vibrating tactors. In vision, mechanisms such as summary statistics are considered to play an even greater role in the perception of complex textures or real world stimuli than for abstract stimuli (Brady et al., 2016), and it would have been expected that previously observed tactile effects (Harris et al., 2001; Kahrimanovic et al., 2009; Roberts & Humphreys, 2010) should have been even more robust here. Furthermore, unlike in Chapter 2, when the non-target items changed, the values within the array could also change, rather than the same values simply changing locations. This would cause increased differences in factors such as summary statistics, which should have further enhanced the potential benefits from leveraging ensemble encoding mechanisms to aid change detection performance.

However, previous evidence of integration across tactile stimulation has been shown to be dependent on the context. For example, such integration was found to be more reliable when the palms were facing downward and the stimuli were therefore presented below the hands (Roberts & Humphreys, 2010). This is the opposite of the setup used in the current experiment and, although previous studies have found such integration effects even with palms facing upwards (Harris et al., 2001), it is possible that such situational cues may have weakened any effects observed in both studies in this chapter.

Beyond simple hand position, a similar consideration might be whether vibrating tactors on different fingers might be more likely to be considered as a ‘group’ than the distinct textures used in this experiment. Within vision, ensemble encoding appears to occur more strongly across items that form a distinct 'group' (Corbett, 2017). There is limited analogous on Gestalt effects in the tactile domain but at least some appear to operate broadly analogously across the two domains (Chang, Nesbitt, & Wilkins, 2007; Gallace & Spence, 2011; Kappers & Bergmann Tiest, 2014). The coherent motion and common onsets/offsets of the stimuli used in this experiment might therefore be expected to induce sufficient such grouping mechanisms. It is worth noting though that studies relating to the effects of such 'common fate', in both vision (Sekuler & Bennett, 2001; Uttal, Spillmann, Stürzel, & Sekuler, 2000) and touch (Bach-y-Rita, 2004) tend to focus on figure-ground distinction, able to detect the presence of coherent changes against either a static or random background. Although ensemble encoding has been shown to be biased by features such as similarity, proximity, connectedness and common region (Corbett, 2017), it is possible that not all grouping mechanisms have a similar effect, and that common fate is not sufficient for between-item encoding. If so, this may also partially account for why such tactors show a greater integration across inputs, given their greater similarity. Future studies may benefit from trying to find textures similar enough that they could be considered more likely to come from the same 'group', but maybe differ along a single key dimension (such as the same material but of different levels of roughness).

A potential methodological issue of both Experiment 4 and 5 is the algorithm for selecting the textures to present on each trial. The only restriction set was that either every non-cued location had a different texture, or all have the same texture, between the encoding and test phases. There was no absolute constraint on any repetitions of stimuli. No card panel had two of the same texture, but given there were only five textures used, there was always at least one repeated texture, and never more than two of any one texture. However, there was no restriction on how many of the textures might have a repetition, and in which locations such repetitions could occur. On some trials, each hand may have had the same textures as each other, perhaps presented in the same spatial positions or mirrored such that the two panels were 'symmetrical'. Such forms of encoding might have been expected to increase the utility of ensemble encoding. For example, if a participant felt that the stimuli were symmetrical during the encoding phase but in the test array, with the non-target locations kept the same, the two hands were no longer symmetrical, the participant would have been able to tell that the target must therefore have changed. Importantly, this encoding and subsequent inference could occur without needing to encode explicitly the individual textures themselves, only the ‘relative’ values between inputs. However, such configuration effects were not observed.

Instead, such grouping effects may have served to reduce the memory loads. For example, if the same textures were presented to each hand in Experiment 4, the participant would only need to encode the textures of one hand, and could then apply this across both, reducing the number of items that need to be encoded from 6 to 3. This ‘chunking’ may have reduced the memory load sufficiently that participants were able to rely more heavily on explicit representations for individual inputs, with reduced utility of group-level representations. The effectiveness of such chunking effects would vary between trials, and trials where the same textures were presented to each hand would have been in the minority, but they may have occurred with sufficient frequency to weaken any group-level effects. One issue with this interpretation though is the generally poor performance across both studies. If such chunking were to play a major role in performance, it would probably have been expected to inflate memory performance and estimates of memory capacity above what was observed here. The algorithm could be adapted to stipulate that only pairs of card panels with no more than one texture repetition is used. However, future research could go further and increase the number of exemplars from only five. The more exemplars are used, the less likely any repetitions within a trial would occur.

Increasing the number of exemplars would also have the added benefit of discouraging artificial encoding strategies such as verbal encoding. Although the texture stimuli might not fall into obvious categories in the same way familiar colour categories might, participants may have applied their own labels such as ‘the really rough one’ or ‘the slightly smoother one’. If so, participants’ memory representations may not have been of the tactile sensation but instead as a verbal list of such ‘labels’. This verbal list would lack the low-level perceptual organisation features such as ensemble encoding that is likely to underpin the configuration effect, and so weaken any potential configuration effects when the test array is presented. Even with more exemplars, participants may still attempt to assign verbal descriptions based on relative between-item differences, but with a sufficient number of similar exemplars such strategies would be less likely as it would become more difficult to sustain distinctive labels for each individual texture, and therefore less likely to be adopted.

One proposed manner by which summary statistics may contribute to the observed configuration effects is that on a trial where the participant is unable to remember reliably the identity of the cued target, they may still be able to infer a change if the non-cued items are presented the same as during the encoding phase and if the participant detects a change in the summary statistics of the array (which could therefore only be caused by a change in the cued target, even if the participant does not remember what the previous identity of the target was). It might therefore be though that when the cued target changed to a different texture, that it may have caused too small change to the summary statistics for the tactile system to detect reliably. This may be particularly true given the textures were relatively similar (all textile ribbons). This would be in contrast to when the colour of an item in a visual array may change from red to green, for example, which may either produce a larger change to the relevant summary statistics, or the visual system may be more sensitive to any such changes than the tactile system. It is also worth considering that previous examples of integration between inputs have often been shown when using only two inputs (Kahrimanovic et al., 2009; Yoshioka, Craig, Beck, & Hsiao, 2011). When using either 4 inputs (Experiment 5) or 6 inputs (Experiment 4), the relative contribution of each input to the overall summary statistic would have a weaker weighting. This would again reduce the magnitude of the change if the cued target changes. It is worth noting that although in the visual domain, increasing the number of items in the array seems to strengthen the reliability of group level integration and summary statistics (Robitaille & Harris, 2011), such findings have focussed on participants ability to calculate explicitly the relevant summary statistic, rather than detect a subtle change from only a single item changing, as is the case here.

The configuration effect is only one potential paradigm to investigate memory representations at different levels in a hierarchical structure. Given the relative lack of research on tactile working memory, the absence of any reliable group level effects in this chapter should not be generalised to indicate that no such mechanisms may occur more generally. Future research would still benefit from including further use of real-world textures rather than abstract tactile stimuli in order better to capture any effects from hierarchical encoding. However, the generally poor performance of the current task should encourage researchers to carefully consider their chosen stimuli. For example, the perceptual discriminability of the exemplars used in this chapter were not tested explicitly and so it remains possible that they may not have been sufficiently discriminable, leading to greater confusion and poorer overall performance. Conversely, choosing textures that share more similarities and so can be ‘grouped’ more naturally may encourage such group-level encoding. Although such changes might inform us of the limits of the generalisability of any such mechanisms, the use of textures that form a grouping would be closer to how textures are often encountered in the real world.

Similar to this, research may wish to employ a setup where the participant's palms face downward with the stimuli underneath, as this has been shown to increase the level of integration (Roberts & Humphreys, 2010). This was not adopted in the studies in this chapter due to pragmatic difficulties. Future research may benefit generally though from finding ways to mechanise or automate the presentation of stimuli, to allow for more consistent presentation pressures and speeds than could be achieved by a human experimenter, even despite the tools employed to help improve standardisation of such features. Although the perception of texture features such as roughness appear largely immune to variations in presentation factors such as speed (Lederman, 1983), this 'constancy' seems less reliable for passive presentation (as in these experiments, where the participant remains passive while the textures are presented to them) than for active feeling (where the participant moves their hands to feel and explore the texture; Yoshioka, Craig, Beck, & Hsiao, 2011). Therefore, if future studies adopt a similar design to the current experiments, with passive presentation, more robust controls against such variability may be necessary.

In addition to the presentation method and bodily posture, further thought may need to be given to the area of stimulation. In both experiments in this chapter, care was taken to avoid stimulating the fingertips themselves, for fear that they may be too sensitive and make the task too trivial. As has been outlined above, it is common in tactile perception research to stimulate various parts of the body other than the fingertips, but these tasks often require simple detection of 'any' stimulation at that given location, rather than subtle identification or discrimination of 'what' the stimulation is. It is possible that the tactile receptors away from the fingertips may have lacked the sensitivity to discriminate reliably between the textures used here. Given the low overall performance, future research should aim to replicate the current design but specifically stimulating the fingertips. This is especially true if some of the changes suggested here are adopted, such as increasing the number of exemplars, which may make the task even harder. Some have also suggested that the limited capacity for tactile working memory is even worse at locations on the hand other than the fingertips (Yoshida, Yamaguchi, Tsutsui, & Wake, 2015). However, it should also be acknowledged that the results of these experiments demonstrate that participants were reliably above guessing and were able to perform relatively reasonably what was designed to be a very difficult task, perhaps due to the additional information that can be utilised from the multi-feature real world stimuli rather than unidimensional vibrating stimulation.

As described in Chapter 2, a final consideration is that of statistical power. Participant numbers were chosen to be similar to or greater than those in previous studies reporting similar effects, those studies may also have lacked power to detect such effects consistently and reliably, and therefore basing sample sizes on them may lead subsequent studies, such as those here, to yield false negatives.

Estimating the reasonably expected effect size is difficult as there are few comparable studies in the previous literature. However, studies reporting integration of inputs across fingers, which is argued here to be a potential mechanism of the configuration effect, exhibit effect sizes in the range of d_z=0.5 to 0.7 (e.g. Kahrimanovic et al., 2009). A calculation using G*Power (Erdfelder et al., 2009) revealed that to detect such effects reliably (with a power of 0.8) would require sample sizes between 15 and 27 participants. Therefore, it might be that the sample sizes used in experiments 4 and 5 (approximately 20) might not have been sensitive enough to detect reliably some of the smaller effects previously reported. However, this potential slight lack of power is unlikely to account for the null results entirely as the actual effect sizes reported in this chapter (experiment 4 d_z=0.2 to 0.4, experiment 5 d_z=0.03 to 0.05) were even smaller than those previously reported, especially experiment 5. It does therefore remain possible that a weaker than expected configuration effect may exist in the tactile domain using a paradigm similar to experiment 4, and this smaller effect should be anticipated and accounted for when designing new studies. However, the substantial drop of effects in experiment 5 instead likely reflect an absence of any such effects, due to the methodological changes (such as fewer fingers sampled per hand).

Although there are key questions that have been identified here, many of the recommended changes would require substantial changes to the physical materials used and the method of their presentation, which are not practical for the immediate series of research. Therefore Chapter 4 will instead adopt a different focus. Chapters 2 and 3 have investigated how information at different levels of the hierarchical structure of perception are encoded and may be utilised to aid performance. Chapters 4 and 5 will investigate how already encoded information may bias the perception of new incoming information and, importantly, whether representations at different hierarchical levels exert different influence.

Chapter 4: Guidance of Visual Attention by Visual Working Memory of Different Feature Types

General Introduction

As described in Chapter 1, evidence has shown that information held in working memory can bias the deployment of attention across new incoming information even on a secondary task irrelevant to the memory contents (Olivers, 2009; Olivers et al., 2006; Soto et al., 2005, 2006). The most common paradigm involves performing a search task during the memory retention interval with the critical conditions being the presence of an irrelevant feature singleton also presented. Such an item may generally impair search performance by distracting attention, but the magnitude of this distraction is greater when the irrelevant feature singleton matches the contents of working memory than when it does not (Olivers et al., 2006). However, these effects seem to occur most reliably with low level features such as colour than more complex feature conjunctions such as shape (Downing & Dodds, 2004). Furthermore, although such guidance does seem to be due to the current contents of working memory specifically (rather than priming or previously encoded but no longer relevant information (Olivers & Eimer, 2011; Olivers et al., 2006), not all memory items seem to exhibit such influences. Instead it seems that only a single active search template held in working memory can bias attention and any additional working memory representations are held as ‘accessory’ items (Olivers et al., 2011; Peters et al., 2009). This has manifested in the findings that such memory guidance occurs reliably only if the information for the secondary task is consistent across trials so it can be held in long term memory, and that only a single item is ‘active’ in working memory (Olivers et al., 2011). That is, although such guidance effects may still occur with more than one item in memory, this seems to be the case only if one of them is clearly marked as the ‘active’ representation by other task demands (van Moorselaar et al., 2014).

Critically, this biasing of attention is not unique only to exact matches between the memory item but instead has been shown to generalise also to features that are similar, but not identical to, the contents of memory (Olivers et al., 2006). However, exactly how the magnitudes of such effects vary as a function of the similarity between memory contents and the feature singleton in the search task is unknown. One factor that may be important to consider is the level of representation within the visual system hierarchy from which such memory guidance effects originate (Hochstein & Ahissar, 2002). For example, holding a square in memory might also lead to positive attentional effects not only for other squares present in a scene but also for cardinal (horizontal and vertical) orientational information in general (such as also that of a + shape), or vice versa with a diamond and x shape. Examinations of such effects might help to establish whether memory guided attention effects are specific only at a specific level of representation (in this example, the square or diamond shape), or whether these effects could spread also to other levels of representation (in this example, cardinal or diagonal orientation information).

It is possible that the magnitude of memory guided attention effects may steadily decrease as the similarity between feature singleton and memory item decreases. However it is also possible that holding one value leads to improvements at a broader 'category' level of representation, leading to similar levels of biasing for other values in a channel or category (so remembering one shade of red will have positive impacts on most other shades of red, with a sharp drop off only as the hue begins to transition to another colour category). The focus of this chapter is therefore how the above effects are affected by the exactness of the memory-match (i.e. the degree of similarity between the to-be-remembered item presented at the start of the trial and the memory matching item presented during the search). Previous experiments have found evidence of memory guided attention from one hue of a colour category (e.g. red) to other different hues of the same colour category (Olivers et al., 2006). However, such effects have also been observed when the memory item was a text based description (such as the word "red") while the feature singleton presented in the search task was a stimulus with that colour (Soto & Humphreys, 2007). A further study found that needing to remember how many items were in a memory array can also bias attention towards the alphanumeric character of that number (Moores & Maxwell, 2008)

Such evidence supports the possibility that memory guided attention may not activate only a very narrow range of values close to the specific value in memory but has the potential to spread to a wide range of other values within a given category. The spread of this range may differ depending on the specificity of the memory representation. For example, if the memory task requires participants to remember a very specific shade of red then this precise memory representation may show a narrower spread to other feature values than if the memory stimulus was at a category level (such as the word "red"). Such differences might reflect different influences from different hierarchical levels of representation within memory.

Although effects have been shown for search items with features that are similar to (but not identical to) the remembered information, the possibility that the magnitude of the effect varies as a function of the similarity has not been tested (although the possibility was noted by Downing, 2000). However, to quantify such differences between stimuli would require a feature that can be expressed along a continuous continuum. This chapter seeks to develop a paradigm to allow this.

The use of a continuous measure may also have implications for any effects when more than one item is held in working memory. Although previous studies using designs where more than one stimulus was held in memory have generally failed to find reliable memory guided attention effects when more than one item was held in memory without one being clearly designated the ‘active’ template (Soto, Hodsoll, Rotshtein, & Humphreys, 2008; van Moorselaar et al., 2014; Woodman et al., 2001), it remains possible that multiple memory items may instead cause coarser and potentially more subtle effects. For example, holding in mind an array of items rather than a single item might not lead to reliable differences in memory allocation when simply comparing a binary distinction of ‘same’ versus ‘different’ between the search and memory feature values. However, the combination of memory items may lead to more subtle differences in the shape of the distribution of effect magnitude as a function of similarity, such that feature values around the group average may show improved attention.

Alternatively, memory guided attention effects may instead be affected by the distribution of values in the array. If enough items in the memory array fall within a certain range of each other, feature values in a similar range might yield memory guided attention effects relative to feature values distinct from any such clusters, despite there being multiple items in memory. Such effects may not have been detected in the existing paradigms either because of the choice of values in the ‘different’ category may not have taken into account such group level effects and so been too weak for effects to be found reliably, or if values in the array were relatively evenly distributed then such effects may not have been observed. When holding two colours in mind, it is unknown whether one might observe a binomial distribution or a unimodal distribution centred somewhere between the two values, and whether the slopes around such peaks are symmetrical or biased in one direction, or how the steepness of the slopes may vary under such conditions. Similarly, it remains possible that if the array of memory items was fairly homogenous with colours towards the ‘warm’ end of the spectrum (such as reds, oranges and yellows) that one may observe differences in the magnitude of memory guided attention when the search task colour singleton was another ‘warm’ colour (regardless of whether it matched one of the ‘colour categories’ in the memory array or not) relative to that of a ‘cool’ colour (towards the blue end of the spectrum).

One major problem is that the effect has primarily been investigated using colour, given its reliability in producing the desired effect. However, it can be difficult to use colour as a continuous variable, especially a one-dimension parameter. A similar issue applies to the use of faces or shapes also. As outlined previously, one of the major distinctions between colour and the other commonly used features such as shape, is that colour is a relatively low-level basic feature dimension and so may be able to guide attention more effectively than the latter, which are often complex conjunctions of features. Therefore, the first aim of the current research is to establish whether the effect can be observed using features that lend themselves more easily to unidimensional linear or circular parameters. Two obvious candidates for such features are angle of orientation (such as of a Gabor grating; Experiment 6) and angle of direction of motion (such as of moving dot patches; Experiment 7). Such features also reduce potential effects of articulation (though some of the above studies did use articulatory suppression; Olivers, 2009) than when using distinct categories such as red/green or square/circle.

If the effect can be shown with the features of orientation and motion direction then additional studies could investigate how such effects differ based on differences in similarity. One possibility is that the size of the effect may remain relatively unchanged until the difference reaches a sufficiently large magnitude and then drop off rapidly (for example, the effect may be similar for all shades of red but rapidly disappear as the colour transitions from red to another colour category). An alternative example may be that the magnitude of the effect is greatest for exact matches and is reduced monotonically as the difference increases, such as with a Gaussian distribution.

Either way, a key follow-up question would then be how the distribution changes shape under different circumstances. For example, when very precise memory representations are required (to differentiate between subtle differences), would the memory-similarity effect become more feature-specific (showing an increased centre-surround effect) or more spread (given the increased amount of attentional effort required).

Reaction times are the primary measure of these experiments, but effects of memory guided attention may also exhibit themselves in other ways. For this reason, accuracies on the search task, and on the memory task, are also considered and analysed.

Experiment 6: Guidance of Visual Attention by Working Memory of Oriented Grating

Introduction

The first experiment aimed to establish whether the phenomenon of memory guided attention can be observed when using angle of orientation (of a sinusoidal grating) as the to-be-remembered feature, rather than colour. Participants were asked to remember the orientation of a memory item for later report, at the same time as completing a search task which involved identifying a target number among three distractor letters. An oriented grating (the feature singleton) could appear in a ring surrounding any of the target or distractor items. Critically, the singleton's angle was either similar or dissimilar to the memory item. It was chosen not to use singletons with identical orientations to the memory item to reduce any incentive to 'refresh' the memory representation by attending to the singleton.

It was hypothesised that, when the singleton was at the target location, responses would be faster and more accurate when the singleton was similar (vs. dissimilar) to the memory item. By contrast, when the singleton was at a distractor location, responses were predicted to be slower and less accurate when the singleton was similar (vs. dissimilar) to memory. It was also anticipated that responses would be faster and more accurate overall when the singleton was at the target (vs. distractor) location, and that response times would be slower and less accurate overall when the singleton was similar (vs. dissimilar) to the memory item.

In terms of memory performance, it was predicted that, when the singleton was at the target location, memory performance would be more accurate when the singleton was similar to memory than dissimilar, with no (or reduced) difference predicted when the singleton was at a distractor location. It was expected that memory performance would be greater overall when the singleton was similar (vs. dissimilar) to memory.

Method

Participants

Twenty-one participants (19 female) participated in the study in return for course credit. Ages ranged from 18 to 24 (M=19, SD=1.41) and all participants reported having normal (or corrected to normal) vision. 20 participants were right-handed. All procedures were reviewed and approved by the Departmental Ethics Committee.

Compared to many previous studies demonstrating robust memory guided attention effects, this participant number was comparable (e.g. 20-24 participants; Hollingworth & Beck, 2016; van Moorselaar et al., 2014b; Zhang, Zhang, Huang, Kong, & Wang, 2011) or larger (e.g. 4-15 participants; Beck, Valerie, Hollingworth, & Luck, Steven, 2012; Houtkamp & Roelfsema, 2006; Olivers, 2008; Soto et al., 2005, 2008, 2006; Soto & Humphreys, 2007; G F Woodman et al., 2001; Geoffrey F Woodman & Luck, 2007).

Stimuli and Materials

Stimuli were presented on an LCD monitor, driven at 60Hz, without correcting for gamma. Responses were made on a keyboard. Experiment control and stimulus presentation were programmed using the Psychophysics Toolbox (Brainard, 1997) for MATLAB (2016b, Natick, MA). All stimuli were presented on a mid-grey background.

Memory and search stimuli consisted of apertures presented with an outer edge diameter of 8° with a Gaussian blur at the outer and inner edges. Every pixel of these annuli (prior to applying the Gaussian filter) was randomly assigned as either a light or dark grey pixel. For the search target neutral stimuli, the probability of each pixel being either light or dark was uniform, resembling white noise. However, for memory stimuli and the search singleton stimulus, the probability of each pixel being assigned as light or dark varied spatially as a sinusoidal function with a spatial frequency of 2.3 cycles per degree, at an angle specified by the task parameter for that trial. This variation of pixel density over space gave the impression that the dots clustered together to form oriented gratings made up of light and dark bands.

Each of these annuli contained either a search target or a nontarget. The search targets were always a dark grey ‘8’ or ‘9’ character. Non-targets were chosen randomly (without replacement) from the capital letters ‘COSGQ’. These characters were chosen to minimise oriented lines. Every search stimulus was surrounded by a 1.43° diameter ring. The ring surrounding the search target was a light grey and those surrounding the non-targets were a dark grey. The fixation point was white, with a 0.29° diameter.

Procedure

Participants were seated in front of the screen and were allowed free viewing. Therefore, although they started approximately 57cm from the monitor, no apparatus was used to ensure constant viewing distance. The experiment was split into 5 blocks, each consisting of 32 trials, summing to 160 in total (40 trials where the singleton and search target were spatially congruent and 120 spatially incongruent trials). Within each of these conditions, the singleton was memory-similar on half of trials, and memory-dissimilar on the other half. Participants were free to take breaks between each block and commenced the next block by pressing the spacebar twice. Prior to the start of the experimental blocks, participants completed one practice block of 16 trials (one trial per each combination of experimental parameters).

Each trial consisted of two tasks, a memory task and search task. The memory task involved being presented with a single item to remember to be tested on later. During the retention interval, participants performed a visual search task. Trial progression is illustrated in Figure 6.

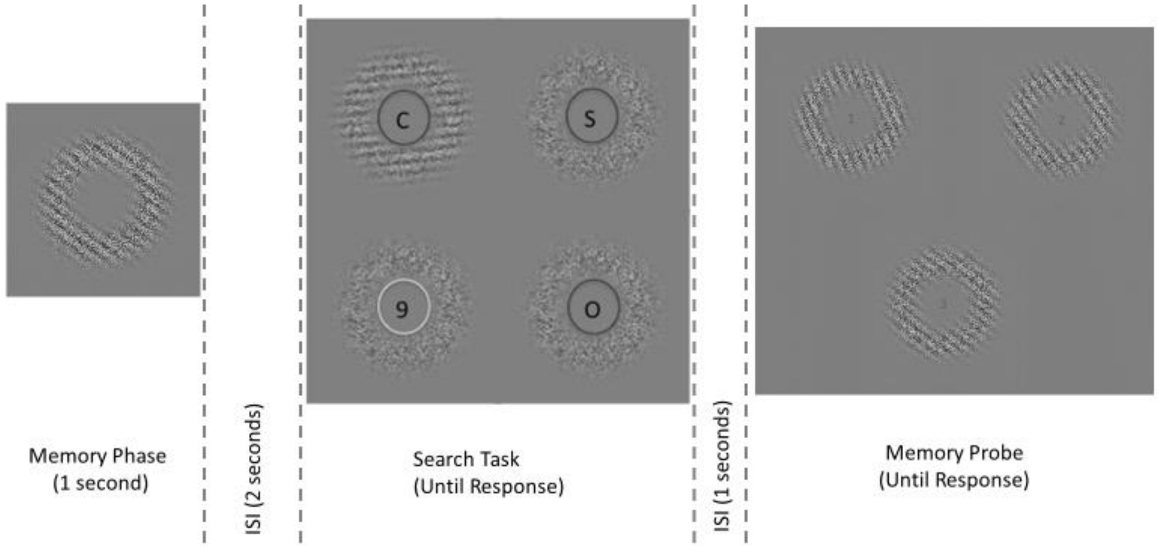

Figure 6: Example stimuli and trial progression. Illustrated trial is of one where the feature singleton appears at a non-target location in the search task and is dissimilar to the memory item.

Each trial started with the fixation-only period of 1 second, to alert the participant to the upcoming stimuli. Following this, the to-be-remembered grating (memory item) was presented at the centre of the screen for a further 1 second. The angle of this grating was determined randomly on each trial.

There was then a 2 second fixation-only retention interval before the search task appeared. The search stimuli consisted of four apertures positioned on the circumference of an invisible circle with a radius of 7 degrees. The angle of each item relative to the true vertical was randomised on each trial, but each aperture was positioned maximally apart from each other, with an arc of 90 degrees between each. Within one of the apertures was the search target (an 8 or a 9, surrounded by a light grey circle) and in the other 3 were distractor stimuli (letters surrounded by dark grey circles). Three of the search apertures contained only white noise and no orientation information, but one (the singleton) contained a grating. This orientation of this singleton could fall either within a range of ±0-15 degrees relative to the memory item's orientation (similar condition) or a range of ±60-75 degrees (dissimilar condition). The singleton had no predictive value of the location of the search target, which had equal probability of being in any of the 4 apertures. Therefore, the target and singleton were in the same spatial location on 25% trials and different locations on 75% trials.

Participants gave their response by pressing either the '8' or '9' key on the 'numpad' using their right hand. For search task responses, participants were instructed to respond as quickly as they could without sacrificing accuracy. The search array remained visible until a response was made and, then immediately removed and followed by another 1 second fixation-period, before the memory test array was presented.

The memory test array consisted of three grating apertures, positioned equidistant (120-degree arc) along the circumference of an invisible circle with a 7 degrees radius, centred on fixation. The position of each item relative to vertical was assigned randomly on every trial. One of these gratings had an orientation that exactly matched that of the original memory item (though the specific luminance values of each pixel were re-calculated independently) and the other two items had orientations pseudo-randomly selected from -30, -15, +15 or +30 degrees relative to the original. If selecting these values truly randomly, the memory-matching item would have been the 'middle' value of the 3 gratings in the probe array on more trials than it would be either of the extremes. To avoid the use of such response strategies, values were constrained such that the memory item had equal probability on each trial of being the far 'left' extreme, middle, or 'right' extreme.

Within each grating was a dark red number '1', '2' or '3', arranged such that each aperture was numbered starting from the left horizontal meridian, moving in a clockwise direction. Participants pressed the corresponding number key in the top left of the keyboard using their left hand. For memory responses, participants were instructed to prioritise accuracy and that there was no advantage to responding quickly. If a response was incorrect, a short feedback beep sounded. Once a response was made, the screen went blank for a 1 second ITI before the onset of the fixation point, indicating the start of the next trial.

Design

Three separate measures were employed in this experiment all within the same 2x2 repeated measures design with the factors of singleton similarity to the memory item (similar or dissimilar orientation) and singleton location (distractor or target). The dependent variables for each analysis were: correct reaction time (RT) to indicate the target identity; target identification accuracy (%); and memory item accuracy (%).

Data Analyses

For all analyses, any trials faster than 300ms or slower than 3000ms were excluded as either anticipatory or abnormal responses (with 0.4% trials removed in this way). Of the remaining trials, any with an RT of more than 2.5 standard deviations from the respective group mean were also excluded (2.3% trials).

For the RT analysis, trials where the search response was incorrect were excluded (3.3% trials) and then trials where the memory response was incorrect were excluded (a further 32.3% trials). All remaining trials were then averaged within their relative conditions.

For the search task accuracy analysis, any trials where the memory response was incorrect were excluded (33.4% trials). For the memory task accuracy analysis, trials with search task errors were included in the analysis.

Results

Search Reaction Times

Mean and standard deviation reaction times are presented in table 12 and figure 7 below.

Table 12: Mean (and standard deviation) reaction times (ms) for search task in each condition

		Singleton Orientation		
		Memory Similar	Memory Dissimilar	Overall
Singleton Location	Search Distractor	1048 (189)	1026 (181)	1037 (182)
	Search Target	999 (200)	1018 (208)	1009 (198)
	Overall	1024 (188)	1022 (189)	

Initially it appears that, when the singleton was at the target location, responses were faster when it was similar to the memory item than dissimilar, with the opposite pattern arising when the singleton was at a distractor location. However, a 2x2 analysis of variance (ANOVA) did not identify a significant interaction between singleton similarity and singleton location ($F(1,20)=2.821$, $p=.109$, $\eta_p^2=.124$).

It also appeared that responses may have been slightly faster when the singleton was at the target location than at a distractor location, but this was not a statistically reliable effect either ($F(1,20)=2.712$, $p=.115$, $\eta_p^2=.119$). There was also no main effect of the singleton's similarity to the memory item ($F(1,20)=0.005$, $p=.943$, $\eta_p^2<.001$).

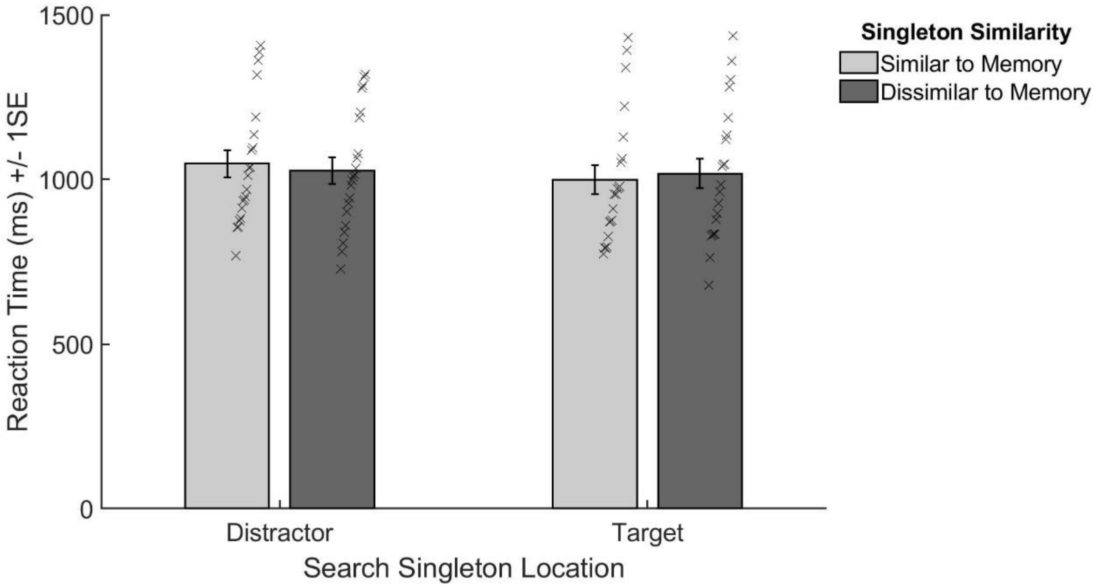

Figure 7: Graph displaying average RTs (+/- 1 standard error) within each condition. Each x represents an individuals' mean RT for that condition (slight horizontal jitter added to display similar results more clearly).

Search Accuracy

Mean and standard deviation accuracy on the search task are presented in Table 13 and Figure 8 below.

Table 13: Mean (and standard deviation) accuracy (%) for search task in each condition

		Singleton Orientation		
		Memory Similar	Memory Dissimilar	Total
Singleton Location	Search Distractor	96.9 (3.7)	96.7 (3.5)	96.8 (2.8)
	Search Target	97.6 (5.4)	94.2 (6.1)	95.9 (4.0)
	Total	97.2 (3.7)	95.5 (3.7)	

Overall responses were highly accurate, but more errors seemed to be made when a singleton at the target location was dissimilar to the memory item than when it was similar (with no such difference when the singleton was at a distractor location). However, no significant interaction was found (*F*(1,20)=1.962, *p*=.177, η_p^2=.089).

There was also no main effect of Singleton Location (*F*(1,20)=1.398, *p*=.251, η_p^2=.065) or Singleton Similarity (*F*(1,20)=3.593, *p*=.073, η_p^2=.152).

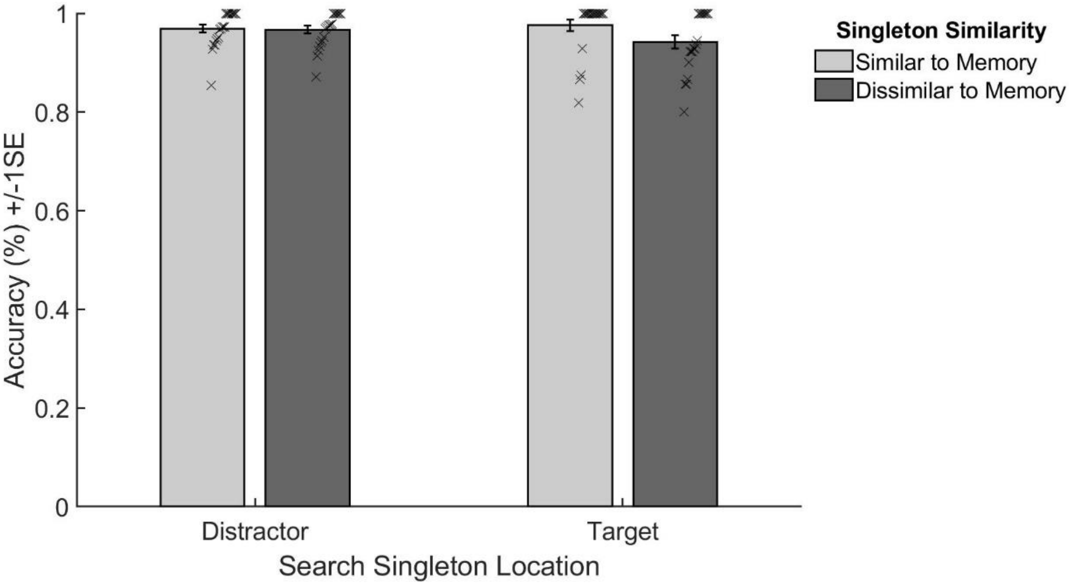

Figure 8: Graph displaying average accuracy (+/- 1 standard error) within each condition. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

Memory Accuracy

Mean and standard deviation accuracy on the memory task are presented in Table 14 and Figure 9 below.

Table 14: Mean (and standard deviation) accuracy (%) for memory task in each condition

		Singleton Orientation		
		Memory Similar	Memory Dissimilar	Total
Singleton Location	Search Distractor	64.1 (9.5)	66.2 (9.2)	65.1 (8.2)
	Search Target	67.3 (12.9)	67.2 (13.0)	67.3 (8.6)
	Total	65.7 (9.2)	66.7 (8.9)	

It initially appears that when the singleton is similar to memory, more errors were made when it appeared at a distractor location than when at the target location (with a reduced difference when the singleton was dissimilar), but this interaction was not statistically reliable ($F(1,20)=0.192$, $p=.666$, $\eta_p^2=.010$).

It also appears as though there were more errors when the singleton appeared at a distractor location than at the target location, but this main effect was also not statistically reliable ($F(1,20)=2.063$, $p=.166$, $\eta_p^2=.094$). There was also no main effect of singleton similarity ($F(1,20)=0.239$, $p=.630$, $\eta_p^2=.012$).

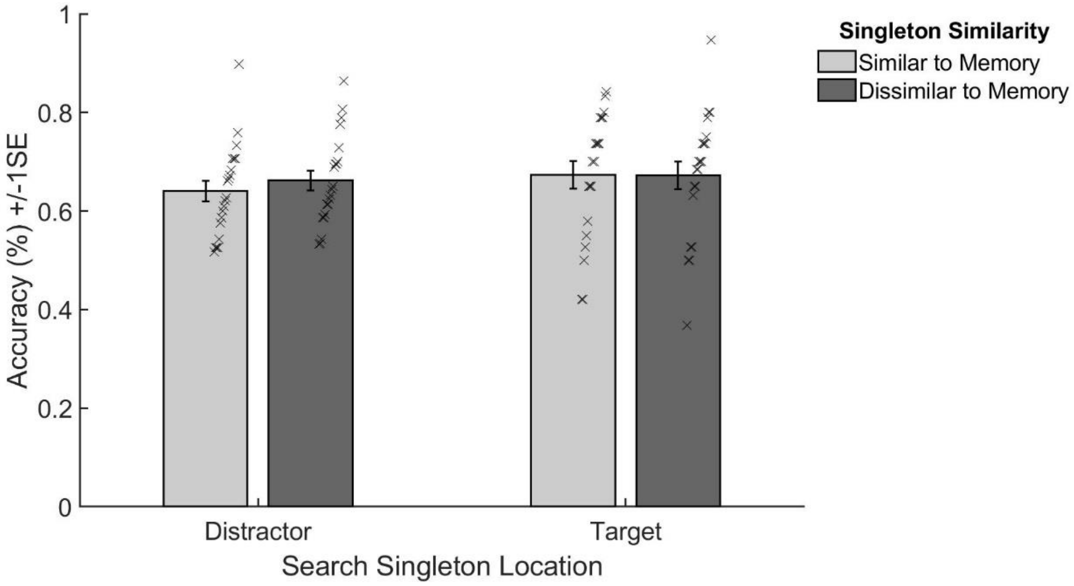

Figure 9: Graph displaying average accuracy (+/- 1 standard error) within each condition. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

Discussion

Qualitatively, the pattern of results appears to be as expected, with a singleton that is similar to the memory item attracting attention more than a singleton that is dissimilar (leading to faster and more accurate responses when at the target location and slower responses when at a distractor location). However, this pattern was not statistically reliable and therefore this experiment has failed to find convincing evidence in support of the hypothesis. However, it is important to remember that this does not mean that this experiment found evidence that there was 'no effect', but that the results are ambiguous and not reliable enough to draw firm conclusions from.

One surprising aspect of these results however is the lack of a reliable difference in response times based on singleton location. It would be expected that a singleton such as the grating would be likely to attract spatial attention due simply to its salience, regardless of whether it matched memory or not (Olivers et al., 2006). Indeed, responses when the singleton was at the target location (biasing attention towards the target) seem to be roughly 28ms faster than responses when the singleton was at a distractor location (biasing attention away from the target), but that this effect was not larger and not statistically reliable is surprising. One possible interpretation of this finding is that the singleton was simply not particularly salient. Efforts were made in the experimental design phase to reduce the saliency of the singleton, such as by introducing a noise mask and reducing the luminance contrast, to avoid a situation in which the attentional capture caused by the singleton was simply too strong to be mediated by the contents of memory. However, perhaps these efforts were too successful and may have been unnecessary given that the singleton in previous studies is often salient (such as a colour among greyscale stimuli; Olivers et al., 2006; Soto et al., 2006).

This may be an issue as the effects of memory guided attention in this type of paradigm are likely to be at least partially dependent on the relative saliencies of the singleton and the target-identifying stimulus (here a white circle among black circles). If the singleton is too salient then it may consistently capture attention (regardless of whether it is similar to memory or not), and if it is not salient enough then it may never be sufficient to bias attention at all (also regardless of whether it is similar to memory or not) or may lead to excess variability from the need to perform a serial search. In previous studies, the target cue was fairly salient and ‘pop-out’ such as a diamond among circle (Olivers et al., 2006), which could be identified efficiently via feature based attention. However, in this study, the target cue was a light grey circle among dark grey circles. Importantly however, these greyscale values are not unique in the scene, with the gratings and random dot apertures having various greyscale values, preventing a true pop-out identification and so again increasing the variability due to the potential need for serial search.

The ideal may be for both the singleton and the target cue to be salient enough to bias attention to a noticeable extent but for the relative competition between them be fairly matched such that attention is either directed directly towards the target or first to one distractor location and then secondly to the target. Considering future experiments, an obvious suggestion is to increase the saliency of both the singleton itself (given that it did not appear to attract attention at all in the current experiment) and the target cue (given that it cannot currently be efficiently identified by feature based pop-out search), both of which will be employed in Experiment 8 later in this chapter.

Another consideration of how to reduce the variability in the search for the target is the number of potential locations. One possibility would be to reduce the number of locations from four to just two. In this way, attention would always shift either directly to the target or first to the distractor location and then to the target. This reduction in possible locations and an absence of a need for a more variable serial search may reduce the noise and variability in the RTs, allowing the qualitative pattern to emerge as more reliable.

However, this approach may fundamentally and qualitatively change the nature of the paradigm and, importantly, be a bigger deviation from established past experiments. Using just two locations would resemble more of a traditional attentional cueing paradigm, whereas previous studies have often used more, not fewer, locations than the four used here (Olivers et al., 2006), though (Downing, 2000 used only two). Although previous studies have found memory guided attention effects even if the search task array was heterogenous and not pop-out (Olivers, 2009), others have failed to show reliable effects (Downing & Dodds, 2004). Such displays may also increase overall variability in response times by the need for less efficient serial searches and may be expected to be less likely to yield reliable results, especially if any present differences are weaker using orientation than those previously observed using colour. Instead, it might therefore make sense to increase the number of locations from four to eight. Doing this may increase the saliency and pop-out nature of the singleton and increase the local contrast (by having stimuli closer together, rather than with such large blank gaps between them). However, it is important that the singleton is kept irrelevant to the search task to avoid the confound of making the feature singleton predictive of the search task target, as in some previous studies (Soto et al., 2005, 2006; Woodman & Luck, 2007) and this requires that the singleton could appear at each location with equal probability. This would require doubling the number of trials therefore (to maintain the same number of trials in the singleton-target location condition), but this may not be desirable, both due to additional resources required and the greater possibility of participant fatigue or disengagement. Therefore, the next studies will continue to use the same four locations used here.

In addition to the reaction time analyses, accuracy on the search task and memory task were also analysed. Performance on the search task was overall very high and ceiling effects and the severe skew of the data mean that firm conclusions should not be drawn from the analysis. However, the key observation is that this analysis did not suggest a speed-accuracy trade off explanation for the RT results. For example, when the singleton was at the target location, performance was both faster and more accurate when it was similar to memory than when it was dissimilar (though again, neither of these differences approached statistical significance). What is interesting though is that when the singleton was dissimilar to the memory item and presented at the target location, performance was slightly worse than when the same dissimilar singleton was at a distractor location. This was not confirmed to be statistically significant so may simply be some random variability. However, should it be replicated in future studies, it is surprising as it would be expected that attention would be drawn more to the target when the singleton was at the target location, even if the orientation was dissimilar to memory, and that this would lead to enhanced processing relative to when the singleton may draw attention to a distractor location. One explanation is that participants may attempt to inhibit the singleton during the search task, as it would be at a distractor location on ¾ trials and so any attentional biasing towards it is more likely to be detrimental to the search task. This inhibition may have different efficacies depending on whether it is similar or dissimilar to the memory item though. A similar memory item may capture attention involuntarily, as predicted, so attempts at inhibition may be impaired, such that it is only when the singleton is dissimilar to memory (so minimal attentional capture occurs) that such inhibition can be observed. Such an explanation may be contradicted by the observed pattern of RT results though, as responses were faster when the dissimilar singleton was at the target location than when it was at a distractor location. Instead, it is possible that the additional perceptual competition from the singleton may have had a masking effect on stimuli at that location, leading to potentially lower accuracy when it was the target at that location. The evidence of the current study is not sufficient to draw any such conclusions though, and the pattern of results would need to be replicated more reliably in future experiments for a more thorough discussion of such effects.

It should also be noted that there was a ‘marginal’ effect of similarity, such that search accuracy may have been slightly higher when the singleton was similar to the memory item, regardless of whether it surrounded the search target or distractor. This may appear surprising, as it would be expected that, if the singleton were able to bias attention, it would only be beneficial if it could bias attention towards the target and away from distractors. Again, the evidence from the current experiment is not sufficient to draw any such conclusions, and such effects would first need to be replicated more reliably in future experiments following improvements to the experimental design. If so, it would be possible that holding information in memory may improve the efficiency of processing similar information, rather than only biasing spatial attention. This could mean that any potential interference or competition from a salient distractor (here the singleton) that is similar to the contents of memory may be reduced because it can be more efficiently processed (and subsequently ignored) than a singleton with dissimilar features, whose processing might be slowed by the competition with the existing internal representation being maintained.

The analyses of the memory accuracy performance also found no significant differences between conditions. Although previous studies have found memory to be better when the feature singleton was at the same location as the search task target than at a distractor location (Moores & Maxwell, 2008), it had been predicted that when the singleton was at the attended target location, mis-integration with the internal memory representation would lead to more relative memory errors when it was dissimilar to memory than when it was similar, with a reduced difference between the memory-match conditions when at a less attended distractor location. Instead, even the qualitative pattern of results did not match this and instead suggest, if anything, that slightly more memory errors may have occurred when the singleton was at a distractor location and especially when it was similar to memory. It is important to emphasise again that this pattern was not statistically significant, so would need to be replicated by subsequent studies before drawing any firm conclusions. However, a putative explanation may be that the mechanism of inhibiting the distractor locations may also lead to an impairment of representations of features at those locations. This would have had little impact on the internal memory representation when those features were dissimilar to the memory item (and may even have been expected to improve memory slightly due to reduced lateral inhibition, though no evidence was found of this), but would lead to greater impairment of the memory representation when it is similar to the inhibited features in the search task.

Overall, the above results do not support the claim that maintaining orientation information in working memory biases subsequent visual attention towards similar orientation information in the environment. Nevertheless, the qualitative pattern, and the identification of several potential flaws of the current design, do give encouragement for further investigation.

Experiment 7: Guidance of Visual Attention by Working Memory of Direction of Motion Information

Introduction

Experiment 7 aimed to repeat Experiment 6 but using direction of motion rather than grating orientation as the to-be-remembered feature (although note that this experiment was run simultaneously to Experiment 6 and therefore includes none of the methodological changes suggested in the discussion of that experiment). Direction of motion was used here as, like orientation, it is a low-level feature, but one that would also lend itself to linear (or circular) parameterisation in potential future studies, as desired. The design and hypotheses remained identical to Experiment 6, with the only difference that the to-be-remembered feature was the direction of motion of drifting random dots within stationary patches, instead of oriented Gabors.

As in Experiment 6, it was predicted that responses in the search task would be faster and more accurate when a singleton with a similar (vs. dissimilar) motion direction appeared at the target location, with the reverse predicted when the singleton occurred at a distractor location (i.e. better performance with dissimilar (vs. similar) singletons). It was also hypothesised that when the singleton was at the target location, subsequent memory performance would be impaired when it was dissimilar (vs. similar) to the memory item.

Method

Participants

Twenty participants (3 males, 17 females) participated in the study in return for course credit. Ages ranged from 18 to 19 (M=18.55, SD=0.51) and all participants reported having normal (or corrected to normal) vision. All participants were right-handed. One participant failed to complete all blocks and so was not considered in any analyses. All procedures were reviewed and approved by the Departmental Ethics Committee.

Design

The design was identical to that of Experiment 6, except that the critical feature for the similarity/dissimilarity condition was the angle of direction of motion of the memory aperture, and the motion singleton in the search array.

Stimuli and Materials

Stimuli and materials were identical to those of Experiment 6, except for the apertures. Within each square matrix (prior to the Gaussian filter), a random 10% of pixels were assigned as white dots. Each dot disappeared and reappeared at another random position 5 times per second (each dot had a random phase). On each frame, all dots also moved in the direction specified for that aperture at a velocity of ~6 degrees per second for motion apertures, but 0 degrees per second for stationary apertures. If a dot reached the edge of the square matrix, it wrapped around to the opposite edge. The square matrices and Gaussian apertures were stationary on the screen, only the positions of the dots within each matrix moved. Figure 10 shows an illustration of an example trial.

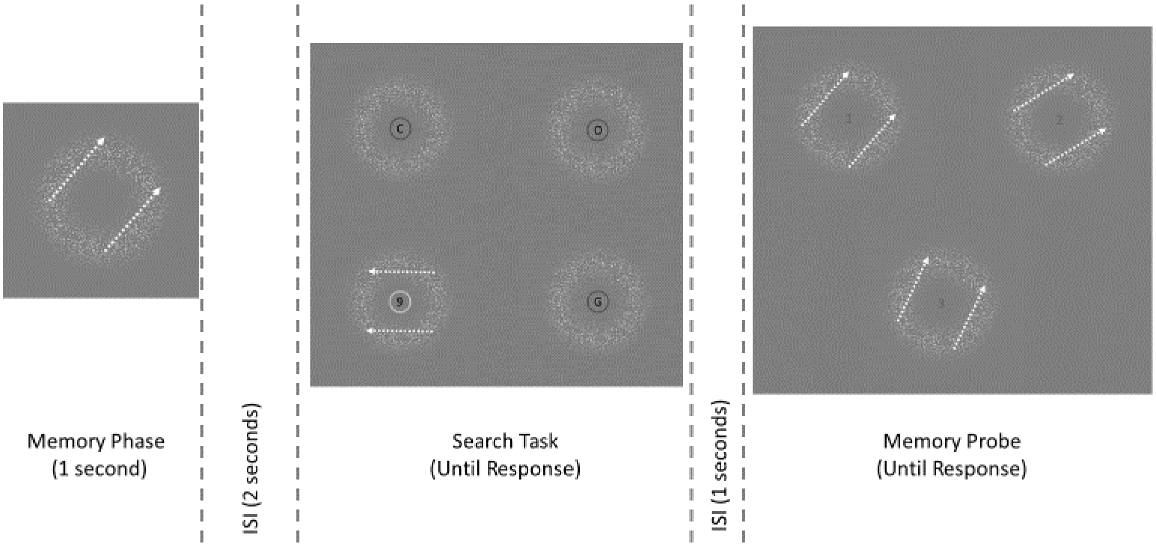

Figure 10: Example stimuli and trial progression. Illustrated trial is of one where feature singleton appears at the target location during search task but is dissimilar to the memory item.

Procedure

The procedure was identical to that of Experiment 6.

Data Analyses

As in Experiment 6, any trials that were faster than 300ms or slower than 3000ms were excluded as either anticipatory or abnormal responses (with 0.8% trials removed in this way). Of the remaining trials, any trials with an RT of more than 2.5 standard deviations from the respective group mean were also excluded (1.8% trials).

For the RT analysis, trials where the search response was incorrect were excluded (2.3% trials) and then trials where the memory response was incorrect were excluded (a further 39.1% trials). All remaining trials were then averaged within their relative conditions.

For the search task accuracy analysis, any trials where the memory response was incorrect were excluded (40.4% trials). For the memory task accuracy analysis, trials with search task errors were included in the analysis.

One participant had reaction times more than 2.5 standard deviations from some group means (specifically the two conditions where the singleton was at a distractor location) and so was excluded from all analyses (final N=18).

RESTULTS

Search Reaction Times

Mean and standard deviation reaction times are presented in Table 15 and Figure 11 below.

Table 15: Mean (and standard deviation) reaction times (ms) for search task in each condition

		Singleton Orientation		
		Memory Similar	Memory Dissimilar	Total
Singleton Location	Search Distractor	1048 (124)	1077 (146)	1062 (132)
	Search Target	982 (170)	955 (120)	969 (129)
	Total	1015 (139)	1016 (118)	

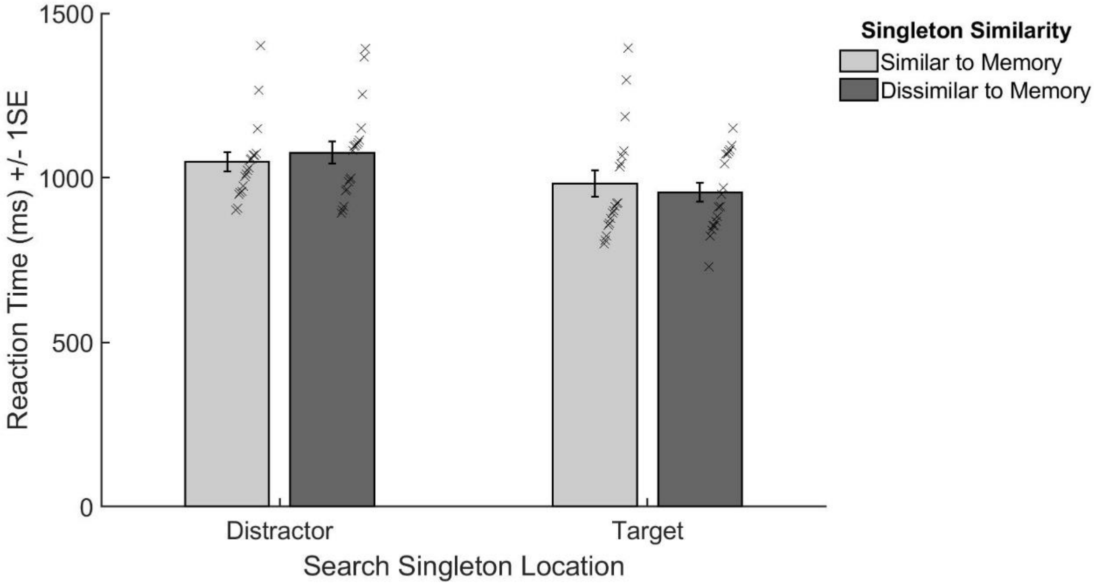

Figure 11: Graph displaying average RTs (+/- 1 standard error) within each condition. Each x represents an individuals' mean RT for that condition (slight horizontal jitter added to display similar results better).

Initially it appears that when the singleton was at the target location, responses were slower when it was similar to the memory item than dissimilar, and the opposite pattern when at a distractor location. However, a 2x2 ANOVA did not identify a significant interaction (*F*(1,17)=2.172, *p*=.159, η_p^2=.113).

It also appeared that responses were faster when the singleton appeared at the search target location than at a distractor location, and this was confirmed statistically (*F*(1,17)=21.076, *p*<.001, η_p^2=.554). However, there was no main effect of singleton similarity (*F*(1,17)=0.002, *p*=.967, η_p^2<.001).

Search Accuracy

Mean and standard deviation accuracy on the search task are presented in Table 16 and Figure 12 below.

Table 16: Mean (and standard deviation) accuracy (%) for search task in each condition

		Singleton Orientation		
		Memory Similar	Memory Dissimilar	Total
Singleton Location	Search Distractor	98.1 (3.0)	97.8 (2.8)	97.9 (2.3)
	Search Target	99.3 (2.9)	98.1 (6.2)	98.7 (3.3)
	Total	98.7 (1.9)	97.9 (3.7)	

Overall responses were highly accurate. No significant interaction was found (*F*(1,17)=0.204, *p*=.657, η_p^2=.012).

There was also no main effect of Singleton Location (*F*(1,17)=1.101, *p*=.309, η_p^2=.061) or Singleton Similarity (*F*(1,17)=0.989, *p*=.334, η_p^2=.055).

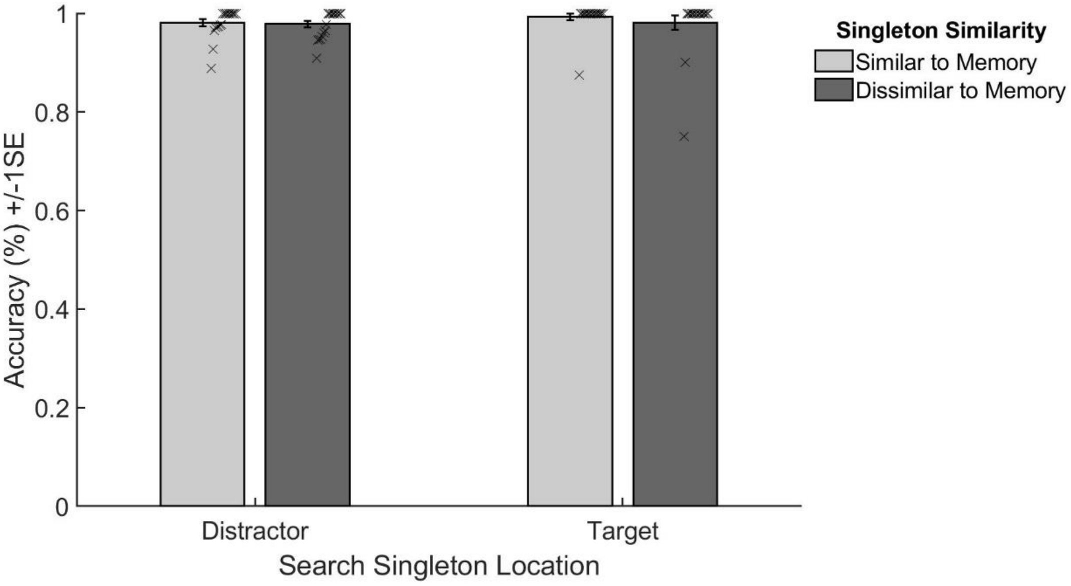

Figure 12: Graph displaying average accuracy (+/- 1 standard error) within each condition. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

Memory Accuracy

Mean and standard deviation accuracy on the memory task are presented in Table 17 and Figure 13 below.

Table 17: Mean (and standard deviation) accuracy (%) for memory task in each condition

		Singleton Orientation		
		Memory Similar	Memory Dissimilar	Total
Singleton Location	Search Distractor	59.7 (10.3)	58.4 (10.8)	59.1 (9.9)
	Search Target	57.6 (10.9)	57.7 (13.6)	57.7 (7.2)
	Total	58.6 (7.8)	58.1 (9.0)	

No significant interaction was found (F(1,17)=0.059, p=.811, η_p^2=.003). There was also found to be no significant main effect of location (F(1,17)=0.375, p=.548, η_p^2=.022) and no significant main effect of similarity (F(1,17)=0.072, p=.792, η_p^2=.004).

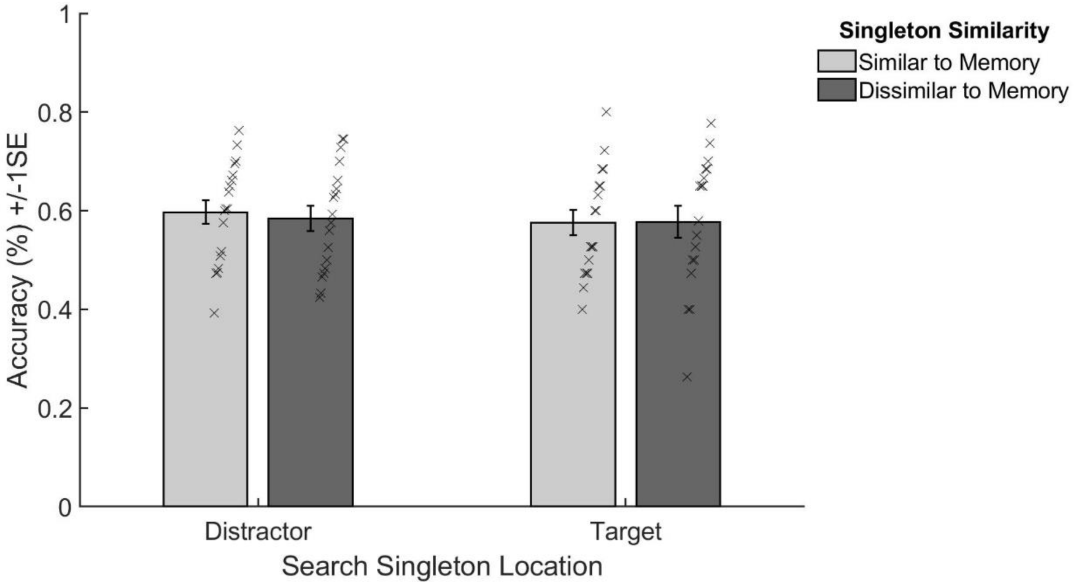

Figure 13: Graph displaying average accuracy (+/- 1 standard error) within each condition. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

Discussion

As in Experiment 6, where orientation was used at the to-be-remembered feature, the results of this experiment found no evidence that holding direction of motion in memory biases spatial attention towards matching motion information in the environment. However, unlike experiment 6, the qualitative pattern of results is the opposite of what was predicted. Here, when the singleton was at the target location, responses were numerically slower when it was like memory, with the reverse pattern of results when at a distractor location. This pattern was not statistically reliable so future studies would need to replicate and confirm whether it is robust, but it would be a surprising finding.

Crucially, and unlike in Experiment 6, response times in the current experiment were significantly faster when the singleton was at the target (vs. distractor) location, suggesting that the singleton in this experiment was sufficiently salient to attract attention in general, but the level of similarity between the singleton and the memory item did not affect the magnitude of this biasing. It is possible that, if the moving dots were too salient, then they may have biased attention so strongly that little opportunity was left for further differentiation based on memory similarity.

Like Experiment 6, search accuracy was consistently high, and no significant differences were observed. This high level of accuracy suggests that participants were engaged and motivated and rules out the possibility that the slight numerical differences observed in the RTs were influenced by a speed accuracy trade off.

Accuracy on the memory task was not affected much by the preceding search task, and no differences were found to be statistically significant.

Overall, these results suggest that the motion singleton did reliably capture spatial attention, but that this did not seem to be affected by the extent to which it was similar to the memory item. When taken together with the results of Experiment 6, these findings suggest that holding orientation information in working memory (in terms of either the orientation of a grating or the direction of motion of drifting dots) does not bias attention to similar incoming sensory information. However, before accepting such a conclusion it was important to replicate the findings of Experiment 6 with a set-up in which attentional capture effects were at least observed overall. The aim of Experiment 8 was therefore to increase the salience of the orientation singleton used in Experiment 6, to ensure that attentional capture effects are elicited in general.

Experiment 8: Refined Design for Guidance of Visual Attention by Working Memory of Oriented Grating

Introduction

Experiment 8 aimed to repeat Experiment 6 (using the orientation of sinusoidal gratings as the memory and singleton stimuli) but to increase the salience of the singleton to ensure that an attentional capture effect is observed overall, whether or not this is modulated by the match between the memory and singleton item. The key differences between Experiments 6 and 8 are that the gratings (both for the memory task and the search singleton) were no longer comprised of clouds of dots with a spatially varying density, but instead were smooth sinusoidal gratings. Also, the target cue was now a green circle (with red circles surrounding the distractors) to allow for a more efficient pop-out search for the target to occur.

The hypotheses are identical to those of Experiment 6 and 7 – specifically, that responses in the search task will be faster and more accurate when the singleton is at the target location and similar (vs. dissimilar) to memory, and the reverse when the target is at a distractor location, where faster and more accurate responding is expected with dissimilar (vs. similar) singletons. It is also hypothesised that, when the singleton is at the target location, subsequent memory performance will be impaired when it is dissimilar (vs. similar) to the memory item.

Method

Participants

Twenty-two participants (7 males, 15 females) participated in the study in return for course credit. Ages ranged from 18 to 28 (M=21.2, SD=2.4), 19 were right-handed, and all participants reported having normal (or corrected to normal) vision. All procedures were reviewed and approved by the Departmental Ethics Committee.

Design

The design was identical to that of Experiment 6.

Stimuli and Materials

Stimuli and materials were identical to those of Experiment 6, except for the following differences. In Experiment 6, a sinusoidal probability function was used to assign randomly the luminance of each pixel, while in this experiment the same sinusoidal function directly determined the luminance of each pixel, varying smoothly between RGB values 77 and 179 (contrast 0.4, though note screen was not gamma corrected). In Experiment 6, the circles surrounding the search stimuli were either a light grey (around target) or dark grey (around distractor), while in this experiment they were either green (target) or red (distractor).

One issue of the more salient sinusoidal gratings is they seemed to create more salient afterimages. To counteract this, a mask was also introduced, and would immediately follow the offset of every aperture in all phases of the trial (including the search phase singleton and non-singletons, and the memory encoding and probe stages) for 500ms. This mask was created by calculating a grid of 5x5 squares, with each square containing a sinusoidal grating, each at a random angle. The boundaries between each square were blurred with a gaussian transition, and the entire grid was itself rotated to a random angle each time it was used. Finally, the outer edges of the grid were removed to leave only a circle of a similar size to the task relevant apertures.

Procedure

The procedure was identical to that of Experiment 6.

Data Analyses

For all analyses, any trials that were faster than 300ms or slower than 3000ms were excluded as either anticipatory or abnormal responses (with 0.5% of trials removed in this way). Of the remaining trials, any trials with an RT of more than 2.5 standard deviations from the respective group mean were also excluded (2.8% trials).

For the RT analysis, trials where the search response was incorrect were excluded (5.2% trials) and then trials where the memory response was incorrect were excluded (a further 33.3% trials). All remaining trials were then averaged within their relative conditions.

For the search task accuracy analysis, any trials where the memory response was incorrect were excluded (35.5% trials). For the memory task accuracy analysis, trials with search task errors were included in the analysis.

Two participants had average reaction time more than 2.5 standard deviations from group means. For one participant this was only when a dissimilar singleton appeared at the target location (though note RTs in all conditions were close to the outlier threshold), and for the other participant RTs were outliers in every condition other than when a dissimilar singleton appeared at the target location (though again, even in this condition RTs were close to the threshold). Therefore, these two participants were excluded from further analyses.

Another participant scored very low on the search task (28%). Chance level would be 50%, and if they had accidently reversed the response-key mappings then, if corrected, their score would still equate to an accuracy of only 72%. This is itself much lower than any other participants across all three experiments, so it is unclear how or why this participant achieved this accuracy. This participant was therefore also excluded from all analyses (final N=19 for all analyses).

Results

Search Reaction Times

Mean and standard deviation reaction times are presented in Table 18 and Figure 14 below.

On inspection of the data there were no apparent numerical trends between any of the conditions. This was reflected also in the statistical analyses, as the interaction was not statistically reliable ($F(1,18)=0.985$, $p=.334$, $\eta_p^2=.052$) and nor were the main effects of distractor location ($F(1,18)=1.161$, $p=.295$, $\eta_p^2=.061$) or singleton similarity ($F(1,18)=0.009$, $p=.927$, $\eta_p^2<.001$).

Table 18: Mean (and standard deviation) reaction times (ms) for search task in each condition

		Singleton Orientation		
		Memory Similar	Memory Dissimilar	Total
Singleton Location	Search Distractor	821 (112)	815 (102)	818 (106)
	Search Target	825 (116)	832 (122)	829 (115)
	Total	823 (113)	824 (107)	

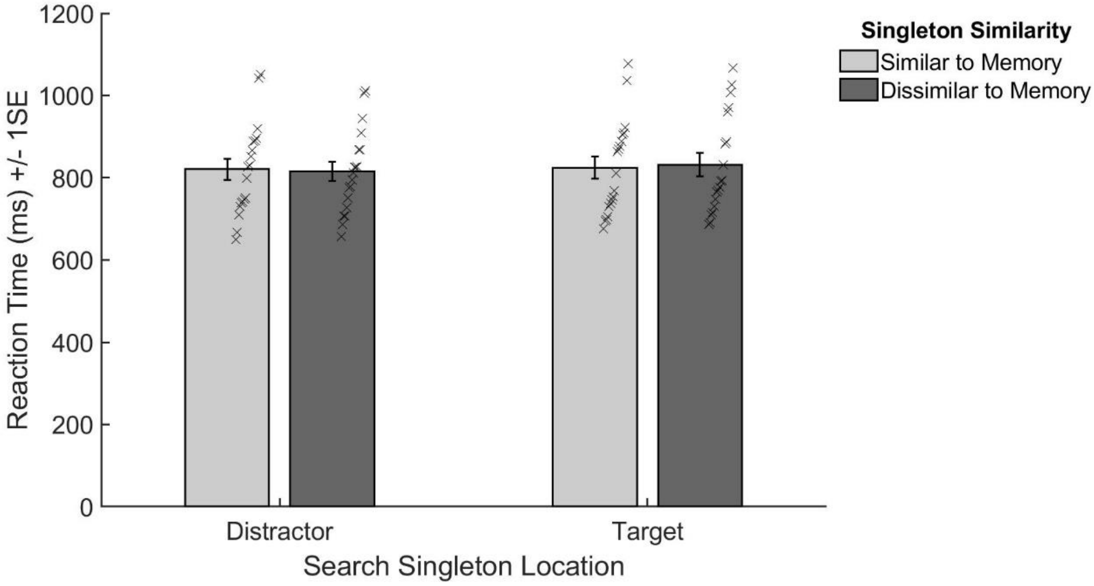

Figure 14: Graph displaying average RTs (+/- 1 standard error) within each condition. Each x represents an individuals' mean RT for that condition (slight horizontal jitter added to display similar results better).

Search Accuracy

Accuracy on the search task is presented in Table 19 and figure 15 below.

Table 19: Mean (and standard deviation) accuracy (%) for search task in each condition

		Singleton Orientation		
		Memory Similar	Memory Dissimilar	Total
Singleton Location	Search Distractor	98.2 (1.9)	98.5 (3.3)	98.3 (2.3)
	Search Target	97.4 (5.1)	96.7 (4.2)	97.1 (3.2)
	Total	97.8 (2.7)	97.6 (3.0)	

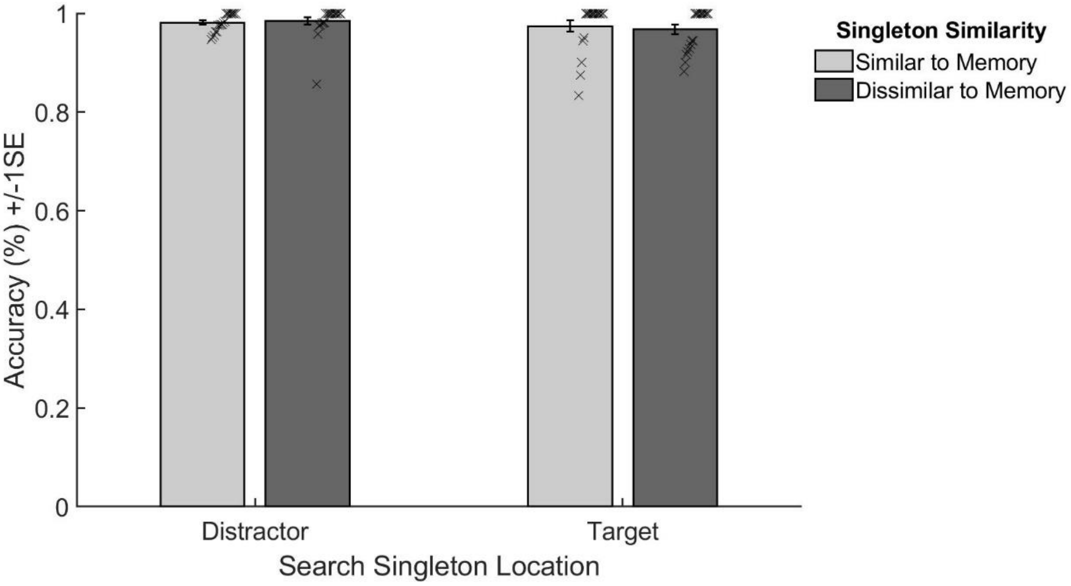

Figure 15: Graph displaying average accuracy (+/- 1 standard error) within each condition. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

Overall responses were highly accurate and no significant interaction was found (*F*(1,18)=0.367, *p*=.552**,** η_p^2=.020).

There appeared to be slightly more errors when the singleton appeared at the target location, but the main effect of Singleton Location was not significant (*F*(1,18)=2.302, *p*=.147**,** η_p^2=.113). There was also no main effect of Singleton Similarity (*F*(1,18)=0.040, *p*=.843**,** η_p^2=.002).

Memory Accuracy

Mean and standard deviation accuracy on the memory task are presented in Table 20 and Figure 16 below.

Table 20: Mean (and standard deviation) accuracy (%) for memory task in each condition

		Singleton Orientation		
		Memory Similar	Memory Dissimilar	Total
Singleton Location	Search Distractor	66.0 (11.0)	65.0 (11.0)	65.5 (10.5)
	Search Target	64.8 (15.3)	57.5 (13.3)	61.1 (12.2)
	Total	65.4 (12.4)	61.2 (11.0)	

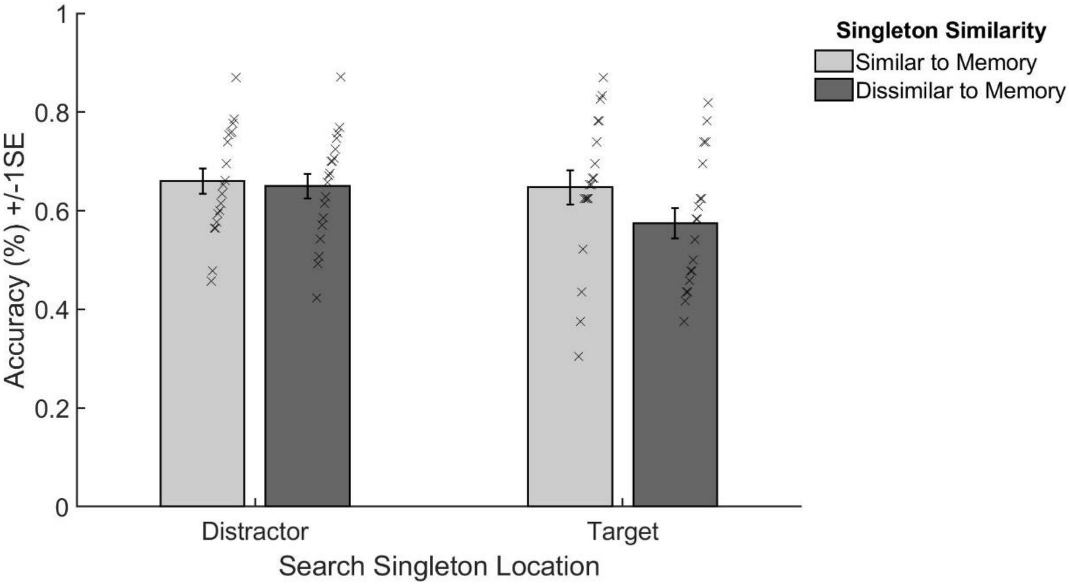

Figure 16: Graph displaying average accuracy (+/- 1 standard error) within each condition. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

When the search singleton is similar to the memory item, it does not appear to have much effect on subsequent memory performance, regardless of whether it appears at a target or distractor location. However, when the singleton is dissimilar to the memory item, subsequent memory performance appears to be noticeably lower when it appears at the target location than when at a distractor location.

This interaction was confirmed statistically (*F*(1,18)=4.840, *p*=.041, η_p^2=.212). Pairwise comparisons confirm that when the search singleton was at a distractor location, there was no significant difference between when the singleton was similar or dissimilar to the memory item (*t*(18)=0.646, *p*=.526, d_z=0.148). However, when the singleton was at the search target's location, subsequent memory performance was significantly lower when it was dissimilar to memory than when similar (*t*(18)=2.110, *p*=.049, d_z=0.484). This interaction also appears to have driven a significant main effect of location (*F*(1,17)=5.369, *p*=.032, η_p^2=.240) and a marginal effect of similarity (*F*(1,17)=3.334, *p*=.085, η_p^2=.164).

Discussion

Like Experiment 6, no differences in RT were found between any of the conditions. Of particular note here is that, again, there was no reliable evidence of any spatial attentional capture at all, regardless of memory similarity, as responses were no faster when the singleton was at the target location than at the distractor location. Unlike in Experiment 6, it seems unlikely that this was due to the singleton not being sufficiently salient, as it was now physically distinct from the random pixel neutral apertures, with a reasonably high contrast. However, in this experiment the target saliency has also been greatly increased (such that it is now indicated by a green circle among red circles), allowing for more efficient pop-out identification. It is possible that this change may have increased the target saliency too much, allowing it to attract attention too efficiently and negating any potential competition from the singleton at all, let alone dependent on similarity to memory. This conclusion would be surprising, as previous studies showing memory guided attention effects have used salient pop-out target cues, such as a diamond among circles (Olivers et al., 2006). However, it may be that in this paradigm of otherwise greyscale stimuli, the use of colour as the target-defining dimension may have been too effective at biasing attention, such that effects from orientation information in memory cannot have a reliable effect (especially given that such information is explicitly irrelevant and even detrimental to the current goals during the search task). Again though, previous studies have used colour singletons as the memory-matching singleton in otherwise greyscale search arrays and still found a difference in the magnitude of memory guided attention between when that singleton matched a colour in memory or not (Olivers et al., 2006). Future studies may wish to use shape rather than colour for this reason. The initial concern with this was to avoid introducing other orientation information as this might then interact with orientation information within memory. However, this may be worth the trade off and, if the target cue is a diamond, it should still pop-out among the circles and circular apertures surrounding each item.

Again, accuracy on the search task was consistently very high, with no significant differences between any conditions. More interestingly, it was found that when the singleton was at the target location, subsequent memory performance was impaired when it was dissimilar (vs. similar) to the memory item, with no such difference when the singleton appeared at a distractor location. This is contrary to both Experiments 6 and 7, neither of which found such difference (not even a similar qualitative pattern). The lack of a significant main effect between when the feature singleton appeared at the search target location or distractor location also contradicts previous studies suggesting performance should be better in the former (Moores & Maxwell, 2008). However, this pattern of results is consistent with what was originally hypothesised and could be accounted for by mis-integration of information at the attended target location. If processing at the target location may also cause some of the surrounding singleton information to be integrated into the internal memory representation, this would lead to reduced accuracy if this erroneous information were more dissimilar to the true memory item than when more similar. Such a difference may not be expected when the singleton is at a distractor location, as minimal processing may occur at this location, reducing the potential competing effects of the singleton on the memory.

It is worth noting that the inclusion of the mask was new relative to both experiments 6 and 7, and also previous research (Olivers et al., 2006). However, it is unclear how this may have affected the results in any meaningful way. It occurred after every stimulus, both in the memory and search task, and every mask was generated randomly and independently each time. Any effects, such as by increasing task difficulty or adding additional orientation information into the paradigm, would be expected to have equal effects across conditions. It seemed likely that it may make the memory task more difficult as it may be expected to interrupt the encoding and consolidation period relative to experiments 6 and 7, but, though there is perhaps a very slight drop in overall accuracy in this experiment, it is minor and still higher than for direction of motion in Experiment 7, suggesting that the null effects cannot simply be explained due to poorer memory representations.

General Discussion

One common implication both from the above studies and previous research (Olivers, 2009), is that effects of memory-guided attentional capture might only occur when specific experimental set-ups are used. If the effects are only reliably observed when colour is used as the to-be-remembered feature, when the search object is well practiced, and when working memory is loaded with only one other feature, then one might ask whether it is really a phenomenon that warrants much further study, whether it is likely have real-life impacts, and whether some useful intervention could be implemented.

Although the current studies might suggest that the phenomenon simply does not occur for orientation angle information, it remains possible that the current parameters are not optimal for revealing any potential effects. It is therefore proposed that future experiments should not model themselves so closely on Olivers et al. (2006) paradigm, but should instead investigate parameters that are maximally likely to induce such an effect. If no effect is observed even then, then this may more convincingly imply that such effects do not exist. By contrast, if such an effect is observed then follow up studies could be performed to investigate the limits of this effect, and which parameters it is most sensitive to, and to control for alternative explanations. As suggested above, one possible change would be the number of stimuli. One recommendation could be to reduce the number of stimuli, such as to only two, as this would help to reduce variability caused by different serial search patterns by ensuring attention shifts and saccades follow only one of two possible patterns. However, memory guided attention effects seem to be more reliably in conditions of high perceptual load (Koshino, 2017), and so increasing the number of items may actually help lead to more robust effects.

Establishing whether alternative measures could reveal memory capture effects in conditions where they have not previously been found would be informative both for our basic understanding of the underlying mechanisms of attention and memory, and for the more practical issues concerning which situations this may have implications for and what might be done to ameliorate them. However, without replicating memory guided attention effects previously established using colour but with an alternative feature such as angle of orientation, that can be quantified as a continuous circular measure, it is not possible to progress to the more advanced questions that formed the original rationale for this series of research (for example, how the magnitude of attention effects vary as a function of the similarity between the memory item and search task feature singleton, whether different hierarchical levels of working memory representation produce different patterns of such spread, and whether holding multiple items in working memory may similarly manifest as a change in shape of such spread). All such questions would require more precise measures of likely more subtle effects and would depend upon the ability to measure and detect such effects reliably, which the current experiments have failed to achieve.

As in the previous chapters, it is worth considering whether alternative encoding strategies, such as verbal labelling, may have played some role in the lack of expected results in this chapter. The stimuli used here (orientation and direction of motion) do not lend themselves as obviously to such labelling as distinct colours, and unlike in chapter 3 with its limited number of exemplars, each stimulus is sampled randomly from a full continuous scale and requires subtle discriminations between similar values (including +/- 15 degrees). Nevertheless, such strategies could play some role, such as by relating stimuli either to quadrants or compass directions (such as “North North East”). Such strategies might be unlikely though as they would yield reliable benefits on only a minority of trials, where random parameters resulted in such a label not also applying to either of the non-targets in the test array.

As in the previous chapters, a final discussion relating to the null results presented in this chapter is of statistical power. Sample sizes were chosen to be similar to or even greater than previous studies reporting such effects, but it remains possible that such previous studies were themselves underpowered and not a reliable benchmark. Although the use of orientation and direction of motion in such a paradigm is novel, studies using other features such as colour have often reported consistent effect sizes in the region of η_p^2=.4 to .55 (Olivers et al., 2006; Soto et al., 2005; van Moorselaar et al., 2014; Woodman et al., 2001; B. Zhang et al., 2011). An analysis using G*Power (Erdfelder et al., 2009) confirmed that a sample size between 10 to 15 participants would be needed to detect effect sizes in a similar range with sufficient reliability (with a power of 0.8). This may account for why many previous studies have used fewer participants and suggests that the experiments reported here should have had sufficient power to detect the expected effect sizes.

However, the relevant actual effect sizes reported in experiments 6 to 8 were much smaller than those reported in the literature using other feature dimensions (at most η_p^2=.124 using orientation in experiments 6 and 8, and η_p^2=.113 using direction of motion in experiment 7) and would have required sample size of approximately 60 to detect reliably. Although the null results reported in this chapter cannot be explained simply as due to insufficient power to detect effect sizes similar to those using colour and other feature dimensions, the current results suggest that any memory guidance effects (if present) are substantially weaker using orientation and direction of motion, and the current sample sizes lacked sensitivity to detect these unexpectedly weaker effects. Future studies investigating this further should anticipate and account for such weaker effect, collecting appropriately larger samples while also attempting to address some of the methodological considerations discussed here to strengthen any potential effects.

Instead of adapting the parameters further, the final empirical chapter in this thesis will instead investigate the other identified open question concerning the underlying mechanism of these effects. Most previous studies have presented the target information and the memory-matching information simultaneously, and it appears that the likely mechanism underlying the RT differences is that spatial attention (and perhaps the first saccade) is biased such that it is more likely to shift spatially to a memory-similar item than a memory-dissimilar item. It remains unclear how such effects, if any, might operate across more dynamic or changing displays. Chapter 5 therefore aims to investigate how the presence of memory matching information in an attentional blink task may affect temporal processing of sequentially presented information.

Chapter 5: Effects of Visual Working Memory on Temporal Visual Attention

General Introduction

Chapter 4 described some of the ways in which visual attention and visual working memory are considered to interact; in particular, how the current contents of working memory, or what is being held 'in mind', can bias the processing of subsequent information in favour of input that is similar to the memorised information relative to dissimilar input (e.g. Olivers, 2009). However, these paradigms have primarily focussed on how memory biases competition between simultaneously presented stimuli. Yet, as outlined in Chapter 1, attention can also operate over time such that, depending on the task demands, the efficiency of processing new incoming information may fluctuate over time. This chapter aims to investigate whether the current contents of working memory might also bias the temporal processing of stimuli that share similar features.

As described in Chapter 1, a common approach to investigating the temporal dynamics of stimulus processing is the attentional blink (AB) paradigm (Raymond et al., 1992). This involves presenting participants with a rapid stream of stimuli (RSVP, typically 10Hz), one at a time at the same spatial location. Embedded within this stream are two 'targets' that need to be remembered and recalled at the end of the trial. Typically, participants have high accuracy for reporting the first target (T1), but accuracy for the second target (T2) depends on the temporal interval between the two. There is marked decrease in T2 accuracy if it presented a few hundred milliseconds after T2, with this period termed the 'attentional blink'.

This chapter employs the AB paradigm to investigate whether the level of similarity between some remembered information and the properties of stimuli in an AB task might affect the magnitude of the apparent blink. One example of such an effect could be when participants search an RSVP stream for even a single target, presenting a memory-similar (vs. dissimilar) item a few hundred milliseconds prior might impair identification of the target (i.e. whether a memory-matching item would induce a 'attentional blink' even with only a single target). Evidence for this possibility comes from a study finding similar results: when participants were asked to conjure a mental image of a type of animal during an RSVP task, presenting an image of that animal (vs. a different animal) created a greater subsequent AB (Pashler & Shiu, 1999). Given the close relationship between imagery and working memory (Tong, 2013) it would seem likely that a working memory version of this task should produce similar effects.

Further evidence for this comes from a similar study where participants were required to remember a word while performing a subsequent RSVP task (Sasin & Nieuwenstein, 2016). A line drawing related to the memorised word appeared to cause an attentional blink of a secondary RSVP target. However, this experiment did not include a condition with an unrelated line drawing, and the temporal positions of the critical stimuli were constant on all trials. These studies will address these issues by including both related and unrelated singletons in the RSVP task, and reducing the predictability of each critical stimulus.

However, previous studies investigating similar questions (Akyurek & Hommel, 2005; Koelewijn, Van Der Burg, Bronkhorst, & Theeuwes, 2008; Peters, Goebel, & Roelfsema, 2009; Dowd et al., personal communication) have failed to find any such effects. When T1 or T2 are identical to items in working memory, they are sometimes less likely to be identified correctly overall, regardless of their serial position (i.e. it is not simply that the blink is amplified, but identification is worse even at non-blink lags; (Akyurek & Hommel, 2005; Nieuwenstein, Johnson, Kanai, & Martens, 2007). The authors suggest this may be due to difficulty in associating the same stimulus to multiple tasks, such that if a letter is already associated with the memory task, it is more difficult for a new simultaneous association to be formed with the RSVP task. Others have found the predicted effects, with a greater blink of T2 when T1 shares features with the contents of working memory (Koelewijn et al., 2008). However, they found similar results when T1 shared features with a similar stimulus simply presented prior to the RSVP stream, even if no memory encoding was required. This might indicate these effects could occur from a form of 'priming', where the mere previous exposure to the memory matching features might affect subsequent processing of similar information, independently of whether it was held in memory or not. This would be different to equivalent effects in the spatial attention paradigms that have established that only items currently held in memory exert such attentional biasing (van Moorselaar et al., 2014). Another study requiring participants simply to identify a target within a sequential stream (not an AB task) while holding another item in memory found no behavioural effects on ability to identify the search target accurately according to whether the memory matching item was also present in the stream or not (Peters et al., 2009). Furthermore, they found that though electrophysiological ERP responses differed for the search target and the non-targets, there was no reliable distinction between such neurological responses to the memory matching item and other non-targets, suggesting its similarity with the contents of memory did not lead to differential processing.

However, there remain some outstanding issues with the designs adopted by some of these studies that may preclude such conclusions and warrant further investigation. In some studies (Peters et al., 2009) the presentation rate was slow (one stimulus per second) and therefore any subtle differences in processing efficiency may not have been detected. Similarly, the search task target changed on each trial (Peters et al., 2009), which has also been shown to reduce any memory guided attention effects in the spatial attention paradigms (Olivers, 2009). In other studies multiple items were maintained in memory simultaneously (Akyurek & Hommel, 2005; Nieuwenstein et al., 2007), which has been found to minimise any memory guided attentional biasing in the spatial attention paradigms (Olivers, 2009; van Moorselaar et al., 2014). The experiments in this chapter will therefore require participants to hold only a single value in working memory, and will keep the search target consistent across trials, to maximise the possibility of observing any memory guided attention.

Furthermore, in some studies (Nieuwenstein et al., 2007), the critical items in the memory and search task were identical as both required letter identification. However, when considering the spatial memory guided attention experiments, often the targets in the memory and search tasks were different. For example, a colour may be the focus of the memory task while a shape was the focus of the search task (Olivers et al., 2006). This approach is important as it allows for a test of memory-guided attentional capture under conditions in which the dimension that defines the memory item (i.e. colour in this example) is different from the dimension that is relevant for the search task (i.e. shape in this example). The current experiments used this approach. In the following studies, colour will be the focus of the memory task but the RSVP task will require the identification of numbers (amongst letter distractors), and the key question will be whether performance on the RSVP task is affected by one of the items having the same or different colour as that in memory.

Much of the existing evidence on the effects of working memory on temporal attention is inconsistent. One possibility is that such effects do exist but are more subtle than their spatial counterparts. Given the similar difficulty in replicating other, potentially more subtle than expected, effects in previous chapters, the experiments in this chapter will be run online in an effort to increase the size of the sample that can be collected and hopefully maximise the possibility of detecting such effects, if they exist. Previous studies have demonstrated that online versions of classic laboratory based designs can lead to similar results (Brand & Bradley, 2012; de Leeuw & Motz, 2016; Hilbig, 2016; Semmelmann & Weigelt, 2017). One common concern with online studies is the precision at which measures such as reaction time can be measured, but the key measure in this chapter will be response accuracy, and the use of distinct colours and alphanumeric characters should remove issues around differences in viewing distances and specific screen settings. Furthermore, eye movements are unlikely to be able to cause systematic differences between the different experimental conditions, as all critical stimuli will be presented at a single central location.

Experiment 9: Effects of Visual Working Memory on Temporal Processing of Sequentially Presented Stimuli

Introduction

As described in chapter 4, many studies investigating memory guided attention have also adopted the approach that the memory target and the feature singleton in the search task should not be an exact match (Olivers et al., 2006), to avoid the potential confound of any behavioural differences being due to participants using the feature singleton to ‘refresh’ their memory performance on the memory task (even if at a small cost to the search task). However, this experiment will purposefully include an exact match to maximise the detection of the contents of memory influencing the processing of stimuli in the RSVP stream. If detected, follow-up studies will be able to explore whether these are due only to such refreshing strategies or not, but the current priority is to establish whether any such effects can be reliably observed under maximally favourable conditions.

Ideally multiple lags would be sampled but each lag tested multiplies the total number of trials required to maintain a reasonable number per condition. Especially given the online nature of the current study, it would be better to keep the total experiment length to a minimum to reduce risks of fatigue, disengagement or drop-out. Therefore, only specific lags will be sampled.

Given these effects may be potentially weak, they may not be detectable (or may not exist) at very early lags (when the blink may be too strong) or very late lags (when the blink may have ceased). Instead, effects from memory are more likely to be detected at mid-range lags, such as if relative recovery from the blink is delayed, or if the magnitude of the subsequent blink is amplified, by T1’s similarity to memory.

Therefore, lag4 will be sampled as it is within the timeframe that is maximally likely to observe any such effects. However, it will still be important to confirm the presence of a blink at all, regardless of any such effects, and so it is also necessary to sample either an early or late lag, distinct enough from lag4 so that a difference would be expected between lags. The current study will sample lag2, when the blink should be at its strongest, as this will allow the overall RSVP streams to be kept shorter, increasing the number of trials that can be included in the given timeframe.

In the current study, it is predicted that overall T2 identification will be worse at lag2 than at lag4, in line with the typical findings of the AB paradigm. It is also predicted that T2 identification will be worse when T1 is of the same colour as the item in memory, than when T1 is of a different colour, due to memory-guided attentional capture. It is also predicted that the differences in T2 identification between the T1-memory similarity conditions will be greater when T2 is at lag4 than at lag2.

Method

Participants

Eighty-four participants (44 male, 39 female, 1 undisclosed) were recruited via an online recruitment platform Prolific (www.prolific.co). This number was because the intention was to include fifty participants in the final analyses and so any participants that did not meet minimum performance threshold (outlined below) were replaced by additional participants prior to full analyses. This number of participants is comparable to or greater than many previous examples of the attentional blink (e.g. 5-18 participants Chun & Potter, 1995; Koelewijn et al., 2008; Nieuwenhuis & de Kleijn, 2011; Nieuwenstein, Chun, Van Der Lubbe, & Hooge, 2005; Olivers & Meeter, 2008; Raymond et al., 1992; or 24-30 participants; Akyurek & Hommel, 2005; Nieuwenstein et al., 2007; Shapiro, Hanslmayr, Enns, & Lleras, 2017), including those showing potential dual task memory effects (12 participants Pashler & Shiu, 1999), and is also greater than those studies demonstrating parity between lab-based and web-based cognitive studies (e.g. 28-50; de Leeuw & Motz, 2016; Hilbig, 2016; Semmelmann & Weigelt, 2017).

Ages were between 18 and 68 (M=29.0, SD=9.7), 9 were left-handed and 75 were right-handed, and all participants reported no long term physical or neurological issues. Three participants indicated that they did not have normal (or corrected to normal) vision and were excluded from the study. The experiment was expected to take approximately 20 minutes and participants were reimbursed £3 for their participation. All procedures were reviewed and approved by the Departmental Ethics Committee.

Materials and Stimuli

General

The experiment was created using the coder interface of Gorilla.sc (Anwyl-Irvine, Massonnié, Flitton, Kirkham, & Evershed, 2018), programmed in a combination of HTML, CSS, JavaScript, and proprietary code from Gorilla itself. Participants were recruited through the online recruitment platform Prolific.co (www.prolific.co). The settings on each were configured such that participants could only complete the experiment on a desktop or laptop computer (not a tablet or phone) and instructions asked participants to avoid environmental distractions while completing the task, such as by turning their phone to silent (though there was no way to ensure this). Similarly, no measurements were made of the display settings of the computer, the speed of their internet connection or general environment. All visual parameters are given here as they were specified in the code, but it is accepted the same parameters may have produced variable displays for different participants, and therefore should be regarded as approximations rather than exact values.

Each trial was split into five key phases, a memory encoding phase, the RSVP phase, the RSVP response phase, the memory response phase, and the feedback screen, with brief blank screen intervals between each. Throughout all phases, the background was set to a mid-grey (rgb(192,192,192). Figure 17 illustrates an example trial.

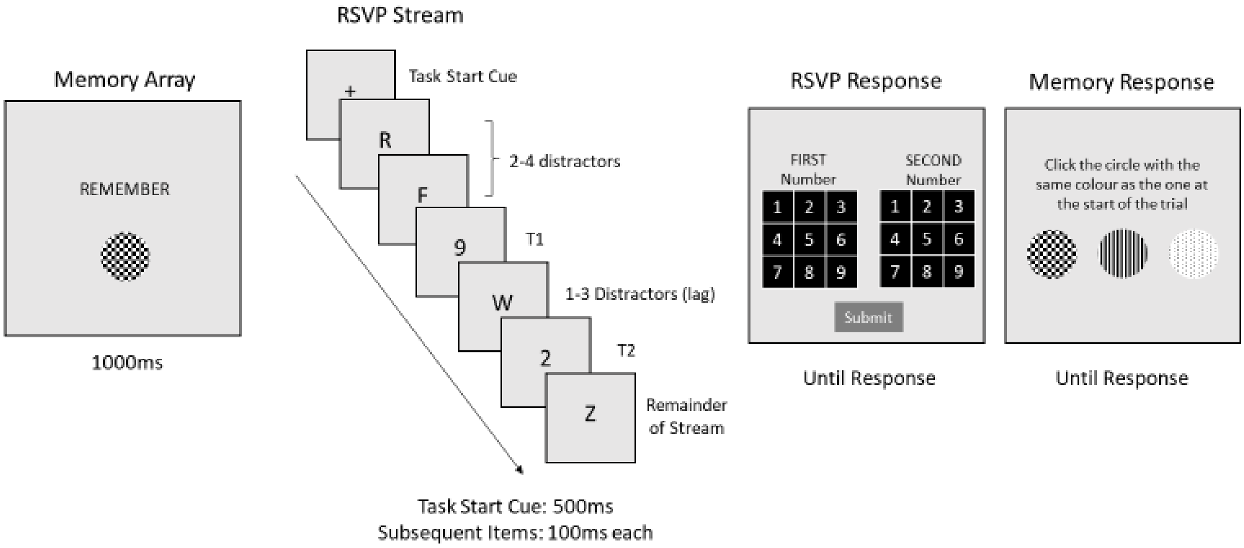

Figure 17: Example stimuli and trial progression.

On each trial some critical stimuli were coloured, described in more detail below. There were fifteen possible colours, split into three categories of green (green phosphor only), yellow (green and red phosphors only) and cyan (green and blue phosphors only). These categories were chosen as on most screens the green phosphor is the brightest, followed by the red phosphor then the blue, and so by using colours with a greater possible luminance range increased the discriminability of adjacent colour exemplars. This was confirmed with some initial piloting which used red, green and blue (instead of yellow, green and cyan), and found that memory performance was best for green and worst for blue.

Within each category were five intensity levels, created by setting the respective phosphors at RGB values of 90, 148, 186, 217 or 243. These values were chosen as an attempt to be roughly evenly spaced intensities between 10, 30, 50, 70 and 90% of each phosphor's maximum intensity, after a default gamma correction of 2.2 was applied. It should be emphasised again though that these values will have varied depending on the monitor, so should not be interpreted as being linearly spaced or as definitely having succeeded in achieving the desired intensities.

Memory Phase

During the memory phase, a single circle was presented at the centre of the screen, with a diameter a little less 20% of the window's height, and the word 'REMEMBER' (font size 30) displayed in black below it. The colour of the circle was chosen randomly out of all the fifteen possible colours and had a black 2-pixel border.

RSVP Stream

During the RSVP stream, a sequence of ten alphanumeric characters (font size 60) were displayed, one at a time, at the centre of the screen. The sequence was preceded by a fixation cross (a '+' symbol) and contained eight letters and two numbers (T1 and T2). The letters were capitalised and chosen randomly from the alphabet (excluding A, B, I, O and S, as these were potentially confusable with the number targets), without replacement. The two number targets, T1 and T2, were each chosen randomly from between 1 and 9, though were always different numbers to each other. Within the sequence of ten characters, T1 could appear in sequential positions 3, 4 or 5 (chosen randomly on each trial) and T2 could appear either 2 (lag2) or 4 (lag4) positions after T1. For example, if T1 appeared at position 3, T2 could appear either at position 5 or 7. The colour of T1 was either the same as the initial memory item (memory same condition) or was chosen randomly from the ten remaining colours that were not in the same colour category as the initial memory item (memory different condition).

RSVP Response

In the RSVP response phase, participants were presented with two virtual number pads, one on the left and one on the right of the screen. These each initially appeared as a 3x3 grid of square black 'buttons', each containing a white number from 1 (top left) to 9 (bottom right). Above the left number pad was the title 'FIRST Number' and above the right number pad was the title 'SECOND Number', both presented in black (both font size 30), to indicate which grid was to be used to indicate the identities of T1 and T2 respectively. Below the grids, in the horizontal centre of the screen was a button. When this phase first appeared, the button stated 'Please Select Both Answers'. Participants were then able to click on any of the buttons to give their responses. When any button was clicked, it changed to a white background with a black number, to highlight that it was the currently selected response. One response could be selected on each number pad at any given time, though participants could change either response by clicking a different button on the respective number pad, which would then highlight in white and the previously selected response for that number pad would revert to the initial black. Once a response had been selected for both number pads, the button at the bottom would change to 'Submit', which participants were required to click to finalise their selected responses (though they were still free to change their responses as much as they wished before clicking the Submit button). If they tried to click this button before a response was given for both numbers (while it still stated "Please Select Both Answers"), it would be ignored and continue to wait for participants to select two responses. Participants were informed to prioritise accuracy and that there was no advantage to responding quickly.

Memory Response

The memory probe phase consisted of three circles presented horizontally, each with a diameter a little less than 20% of the window's height. One was the same colour as the original memory item at the start of the trial and the other two had two random colours chosen from the remaining four intensities within the same category. All three had a black two-pixel border, and beneath the probes were reminder instructions stating, "Click the circle with the same colour as the one at the start of the trial". Participants were informed to prioritise accuracy and that there was no advantage to responding quickly.

Feedback

The feedback screen consisted simply of three lines of text, stating “FIRST number was “, “SECOND number was “, and “Colour memory was “. Each line then ended with either the word ‘Correct’ or ‘Incorrect’ depending on whether the respective response had been correct or not. Underneath the feedback was a button to ‘Start Next Trial’ and would also state which trial number they are on and how many trials there are in total.

Practice Block

At the start of each session was a short 5 trial practice block. Practice trials were identical to the main experimental trials, except for the duration that each character was displayed for (and therefore also the corresponding rate of the RSVP stream), which gradually increased from 400ms (2.5Hz), 300ms (3.33Hz), 200ms (5Hz) then the final two trials at 100ms (10Hz).

Procedure

Participants were first shown a consent screen which informed them of the basic nature of the task and, if consent was given, asked for some basic demographic information. They were then given three screens explaining the task in more detail and then completed a short practice block. They were then shown a reminder of the task instructions (in case anything had not made sense) and asked to confirm that they still wished to proceed with the experiment before the experimental block began. The experimental block consisted of 80 trials (20 trials per four experimental conditions). Trials were self-paced, such that each trial was initiated only when the participant clicked a button to do so.

Each trial started with a blank screen inter-trial-interval of 250ms (except the very first trial, which was preceded by a 500ms blank interval) before the memory phase was presented for 1000ms. Immediately following the memory phase was the RSVP phase, where the RSVP stream was presented. Every stream started with a ‘+’ symbol being presented in the centre of the screen for 500ms, then ten characters for that trial were presented sequentially at the same location, each on for 100ms. Once the RSVP stream had finished being presented, a blank screen was displayed for an interval of 500ms, then the RSVP response phase was displayed. It was decided not to have blank ISIs between each character as preliminary testing on the experimenter seemed to indicate that such blanks made the task harder. There was a desire to avoid the task seeming too difficult to members of the public, to reduce excluded trials and also avoid goal neglect, where participants may disengage and make less effort if they think the task is too difficult anyway (Duncan, Emslie, Williams, Johnson & Freer, 1996).

Participants were required to select one response for the first target and one response for the second target in the RSVP stream, then click the ‘Submit’ button to confirm. Participants had unlimited time to make their decision and could change their responses as much as they wished before confirming with the Submit button. Participants were permitted to choose the same response for both targets if they wished (the same number was never actually used for both targets but was permitted in case a participant remembered one of the values but not whether it was of the first or second target).

Once both responses were confirmed with the Submit button, there was another 250ms blank interval before the memory response phase was presented. Participants were required to select which of the probe items matched the original memory item by clicking on their choice. They were given unlimited time to make their choice but were not permitted to change their response; once they had clicked any of the probe items, this was taken as their response and another 250ms blank interval preceded the feedback screen. The feedback screen was also an opportunity to take a break as the next trial would not begin until the ‘Begin next trial’ button was pressed, for which they were allowed unlimited time.

Once all 80 trials were completed, participants were shown a debrief screen, including information about the study and also a box where they could leave any comments about any issues they may have experienced.

Design

A 2x2 design was employed to investigate the primary question of whether the similarity of a target’s colour with a colour held in memory affects its temporal processing, as evidenced by a change in the magnitude of any subsequent attentional blink. The two independent variables were colour similarity (same or different) and T2 lag (lag 2 or lag 4). These four conditions were counterbalanced such that there were equal numbers of trials (20) with each combination of conditions. The dependent variable was the percentage of trials on which T2 was correctly identified.

Data Analysis

Given the many possible combinations and interactions between the two tasks, and within the RSVP task between the two responses, there were several analyses that could be run on the data to get a comprehensive understanding of any potential differences between the different experimental conditions. Rationale for each is provided in the Results section but the analysis steps were similar for each of them.

First, trials were excluded based on whether the responses for one (or more) of the three responses (two RSVP responses and one memory response) were correct or not. Once these trials were excluded, participants would have different numbers of remaining valid trials per condition. For some of the analyses, where the trial exclusion criteria were strict, many participants had very few valid trials in one or more conditions. Therefore, an absolute minimum performance was set such that any participants who had fewer than eight valid trials in any one experimental condition, or who had a combined average across all four trials of less than 10 valid trials per condition, were excluded. Following the above trial based and participant-based exclusions, participants were then also excluded if any of their group means were more than 2.5 standard deviations from the respective group mean (and therefore considered outliers).

Finally, any remaining trials and participants were submitted to a 2x2 repeated measures ANOVA, with the factors of T1's memory similarity (same colour or different) and T2's lag (lag2 or lag4). The dependent variable varies for each analysis but is always the percentage of the remaining 'valid' trials (after the above exclusions) on which a correct response was given, for one of the three responses (T1, T2 or memory item).

For most of the analyses that focussed on either of the RSVP responses (T1 or T2), there were two analyses run. The first used a 'strict criterion' where a response for T1 or T2 was considered correct only if it was given for the correct target (so if T1's number was given as a response to T2, that trial was considered incorrect). However, participants might occasionally have remembered one or both numbers correctly but not necessarily the correct order, perhaps giving T1's value as the response for T2, or vice versa. Given that the relevant feature for the current study is the identity of the numbers, not necessarily their serial positions, these data were reanalysed with a more 'lenient' inclusion criterion such that a T1 or T2 response was classified as correct if it matched either T1's or T2's value (so for example, if T1's value was given as the response to T2, and an incorrect number actually given as T1's response, the 'T1 response' for that trial was actually considered correct under this lenient criterion).

Results

T2 performance when T1 and memory are correct

The central analysis was whether any processing differences of T1 due to its level of similarity to memory affects subsequent identification of T2, and whether any such effects are sensitive to the temporal delay between the two targets. Therefore, it was necessary to focus on trials in which one can be reasonably confident that both T1 and the memory item were correctly processed.

Strict Criterion: any trials where the T1 response was incorrect (15.4% of trials) or the memory response was incorrect (a further 28.2% trials) were excluded. Following these trial exclusions, 30 participants were excluded as they did not meet the minimum trial thresholds and a further two were excluded as outliers, leaving a final N of 49. The percentage of remaining trials on which T2 was correctly identified were then analysed across the four conditions.

Table 21: Means (and standard deviations) of percentage of correct T2 responses of valid trials following outlined exclusions.

		T2 Lag		
		2	4	Overall
T1 Memory Similarity	Different	80.0 (17.8)	80.1 (17.4)	80.0 (15.2)
	Same	80.2 (18.0)	82.3 (15.6)	81.4 (14.7)
	Overall	80.1 (16.1)	81.5 (15.4)	

There was no statistically reliable interaction ($F(1,48)=0.545, p=.464, \eta_p^2=.011$) and also no statistically significant main effect of either T1 memory similarity ($F(1,48)=0.933, p=.339, \eta_p^2=.019$) or T2 lag ($F(1,48)=0.510, p=.479, \eta_p^2=.011$).

Liberal Criterion: With the extra trials included, only 23 participants did not meet the minimum trial criteria, and a further two were excluded as outliers, leaving 54 participants. This actually led to very slight decreases in accuracy, but this had no effect on the qualitative results. There was no significant difference between lag2 and lag4 ($F(1,53)=0.199$, $p=.658$, $\eta_p^2=.004$), no difference between T1-memory colour similarity conditions ($F(1,53)=0.123$, $p=.727$, $\eta_p^2=.002$), and no significant interaction ($F(1,53)=0.160$, $p=.691$, $\eta_p^2=.003$). This qualitative pattern of results was also true even if only the 49 above participants included in the main analysis were reanalysed this way (T1-Memory colour similarity $F(1,48)=0.446$,

p=.508, η_p^2=.009; lag F(1,48)=0.323, p=.572, η_p^2=.007; interaction F(1,48)=0.758, p=.388, η_p^2=.016).

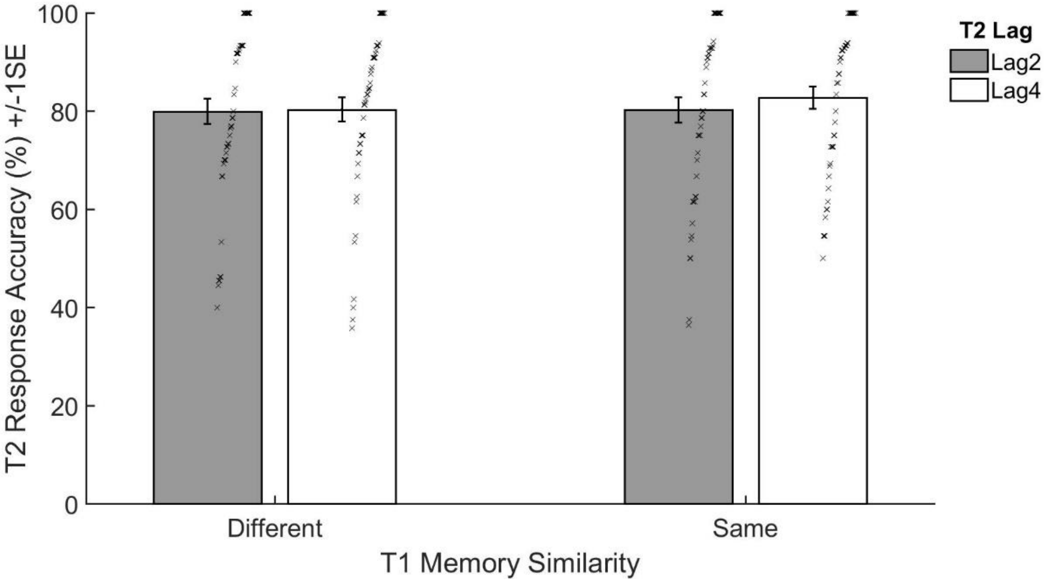

Figure 18: Graph showing mean T2 response accuracy, expressed as a percentage of the number of valid trials after excluding any trials where either the T1 or memory response was incorrect. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

T2 when T1 is correct (disregarding memory performance)

The addition of a concurrent memory task is a relatively new addition to the attentional blink paradigm, so it is worth also analysing the RSVP task in the same manner as previous papers, not considering the memory task, to aid in more direct comparisons. A total of 54 participants were included in the strict criterion analysis and 76 included with the liberal criterion.

With the strict analysis, there was no significant effect of T1-memory colour similarity (Similarity F(1,53)=0.123,p=.727, η_p^2=.002), no significant effect of lag (F(1,53)=0.199, p=.658, η_p^2=.004) and no significant interaction (F(1,53)=0.160, p=.691, η_p^2=.003). Similar results emerged using the more liberal criterion; there was still no significant effect of similarity (F(1,53)=0.099, p=.754, η_p^2=.002), lag F(1,53)=0.866, p=.356, η_p^2=.016) and no interaction (F(1,53)<.001, p=.984, η_p^2<.001).

T1 when memory is correct

Another secondary analysis is how T1 identification may be affected by its similarity (or not) to memory. For example, whether the processing of T1 was affected by working memory, even if the subsequent blink was not affected. Also, errors seemed high generally, and T2 may have interacted with T1, so further analysis of T1 may yield further insights into any influences from working memory.

This analysis excluded trials on which the response on the memory task was incorrect (33.6%), as to include trials where the memory was incorrect may make it differ to infer whether any differences (or lack of differences) were linked to the contents of memory itself. Using the strict T1/T2 criterion, 7 participants were excluded for not reaching the minimum valid trial thresholds and a further 4 were excluded as outliers, leaving a final N of 70 participants. In the liberal analysis, a further two participants were excluded as outliers, leaving 68.

Table 22: Mean (and standard deviations) percentages of trials where T1 was correctly identified, considering only trials on which the memory item was remembered correctly.

		T2 Lag		
		2	4	Overall
T1 Memory Similarity	Different	84.8 (13.3)	88.8 (10.5)	86.8 (10.8)
	Same	84.2 (12.3)	90.1 (10.5)	87.1 (9.6)
	Overall	84.5 (11.4)	89.5 (9.1)	

Strict Criterion: On trials where the memory item was correctly remembered, T1 identification appears to have been greater when T2 occurred at lag4 than when at lag2 (*F*(1,69)=22.171, *p*<.001, η_p^2=.243). There was no effect of whether T1's colour was the same or different to the memory item (*F*(1,69)=0.150, *p*=.699, η_p^2=.002) and no interaction between the memory similarity and T2's lag (*F*(1,69)=1.170, *p*=.283, η_p^2=.017).

This analysis was also rerun including only the same 49 participants included in the core T2 analysis earlier. The pattern of results remained unchanged.

Liberal Criterion: T1 performance was still significantly worse when T2 was at lag2 (M=90.1, SD=8.7) than when at lag4 (M=91.9, SD=7.0; *F*(1,67)=4.114, *p*=.046, η_p^2=.058), There was still no significant difference between T1-Memory colour similarity conditions (*F*(1,67)=0.093, *p*=.761, η_p^2=.001) and still no interaction *F*(1,67)=2.065, *p*=.155, η_p^2=.030).

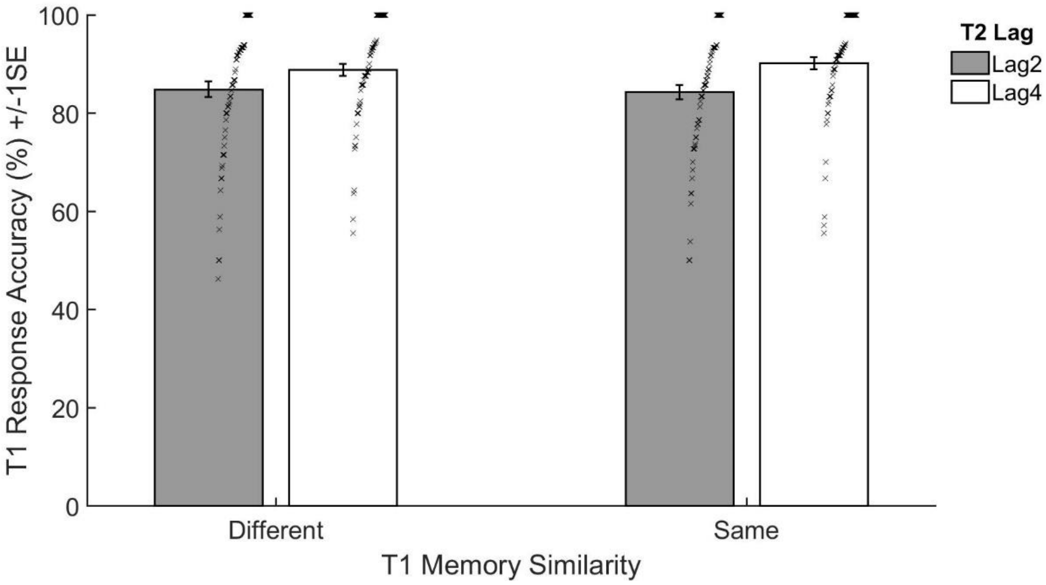

Figure 19: Graph showing performance (% correct) when identifying T1's identity, for only trials where the later memory response was correct, split by T1's similarity to memory and by T2's lag position. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

Overall Data

Although T2 given correct T1 and memory responses was the core analysis for this study, the unexpectedly poor performance and high levels of trial and participant exclusion might have interfered with otherwise present effects. Although this will be addressed in the following experiment in this chapter, to aid with interpretation of the current findings the data were also analysed without excluding any trials based on incorrect responses. Some previous studies have actually adopted a similar approach and included incorrect memory trials in their analyses of search task reaction times (Olivers et al., 2006), though this approach was not adopted in the main analyses as on trials where the memory task was incorrect, there is a higher chance that the memory representation may be incorrect and may no longer sufficiently match the feature singleton in the search task, minimising any such memory guidance effects. Even if such differences were to be observed on such trials, it would then still be problematic to conclude that such effects truly were due to guidance from the current contents of working memory. However, the approach is included here as a secondary analysis to provide further insight when interpreting the core analyses.

Analysis of the memory task is unaffected by whether the strict or liberal criterion is adopted for classification of T1 and T2 responses. Memory performance was significantly worse when T1's colour was the same as the memory item than when it was different *F*(1,73)=25.685,

$p<.001$, $\eta_p^2=.260$). Memory performance did not differ reliably between when T2 was at lag2 or lag4, $F(1,73)=0.585$, $p=.447$, $\eta_p^2=.008$, nor was there an interaction between memory similarity and T2 lag, $F(1,73)=1.833$, $p=.180$, $\eta_p^2=.024$.

Strict Criterion: Seven participants were excluded as outliers, leaving 74 participants. Table 23 and Figure 20 display descriptive statistics for overall performance.

Table 23: Table showing overall means (and standard deviations) when including all trials for percentage correct T1, T2 and memory responses

	Memory Different T2 Lag2	Memory Different T2 Lag4	Memory Same T2 Lag2	Memory Same T2 Lag4
T1 Correct	85.4 (11.4)	88.6 (10.8)	84.3 (11.8)	90.5 (9.2)
T2 Correct	68.4 (21.3)	74.9 (19.3)	68.0 (21.0)	74.8 (20.4)
Memory Correct	70.2 (10.3)	71.0 (14.3)	65.3 (13.6)	62.6 (14.0)

Both T1 and T2 accuracies were statistically better when T2 was at lag4 than lag2 (T1: $F(1,73)=22.012$,$p<.001$, $\eta_p^2=.232$; T2: $F(1,73)=22.333$, $p<.001$, $\eta_p^2=.234$). There were no apparent differences in either T1 or T2 performance when T1's colour was the same as the memory item or different (T1: $F(1,73)=0.190$,$p=.664$, $\eta_p^2=.003$; T2: $F(1,73)=0.046$, $p=.831$, $\eta_p^2=.001$), and neither showed an interaction between memory similarity and lag (T1: $F(1,73)=3.702$, $p=.058$, $\eta_p^2=.048$; T2: $F(1,73)=0.015$, $p=.901$, $\eta_p^2<.001$), though note T1 might be considered 'marginal' and is driven primarily by a slightly greater difference between lag2 and lag4 when T1 was the same colour as in memory than different.

Liberal Criterion: The same participants were excluded as in the strict criterion, but the extra trials included were not uniform across conditions, and the above differences between lag conditions (and any potential interactions with the lag conditions) are weaker in this analysis than in the previous one.

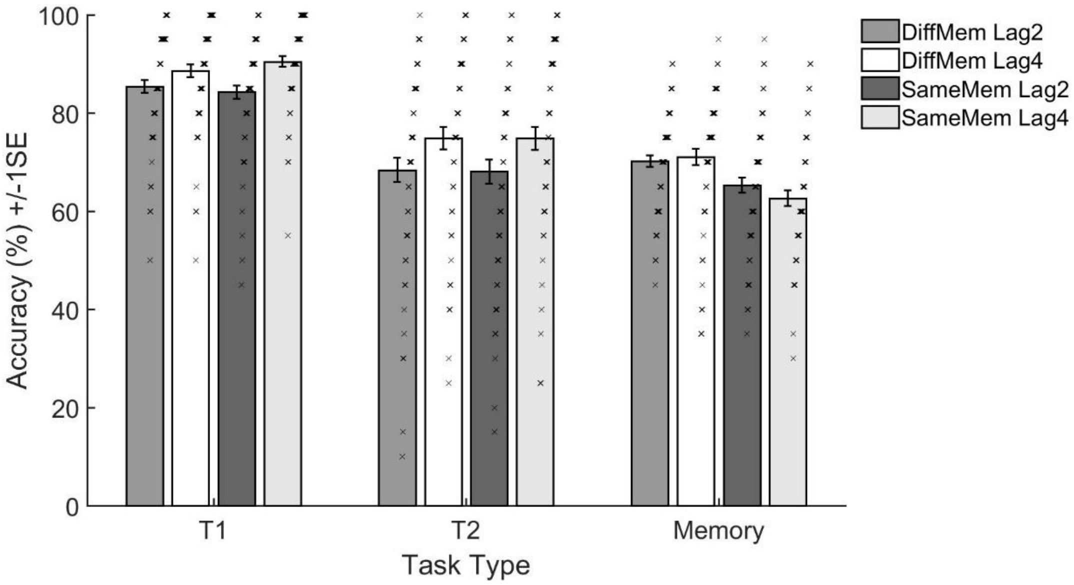

Figure 20: Graph showing overall descriptive statistics when including all trials for all participants. Each small 'x' represents the score of an individual participant. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

For T1 performance, participants were no longer significantly worse when T2 was at lag 2 (M=90.7, SD=7.2) than at lag4 (M=91.7, SD=7.0; *F*(1,73)=1.715, *p*=.194, η_p^2=.023), and the interaction was no longer even marginal (*F*(1,73)=2.055, *p*=.156, η_p^2=.027). There was also still no significant difference between T1-Memory colour similarity conditions (*F*(1,73)=0.209, *p*=.649, η_p^2=.003). Similarly, for T2 performance, performance was no longer significantly worse when T2 was at lag2 (M=76.2, SD=16.7) than at lag4 (M=76.3, SD=17.3; *F*(1,76)=0.010, *p*=.920, η_p^2<.001), and there was still no significant difference between T1-Memory colour similarity (*F*(1,76)=0.086, *p*=.770, η_p^2=.001) and no interaction (*F*(1,76)=0.098, *p*=.755, η_p^2=.001).

Performance Over Time

A final secondary analysis was whether the pattern of performance varied across the duration of the experiment, such as whether performance increased with practice or dropped with fatigue. Given the relatively poor performance and high exclusion rate due to insufficient trials, it is important to understand better how performance may have changed over the course of the experiment. For example, whether performance improves over time once participants have familiarised themselves with an otherwise challenging paradigm and so extending the session to include additional trials may be worth the additional resource investment, or whether performance decreases due to fatigue or waning interest and any such additional trials would only make overall performance even worse.

To this end two analyses were run, one was a series of correlations each between trial number and the average performance of either T1, T2 or memory, and the other was the above 2x2 ANOVA analysis but run on the first half and second half of trials, respectively. Note that for both analyses, all trials and participants were included.

It was found that performance in T1 ($r=.512$) and T2 ($r=.638$) significantly increased with increasing trial numbers ($p<.001$). However, memory performance did not reliably vary over time ($r=-.186$, $p=.099$). It should be noted that although significant, these increases remained relatively small, equating to overall increases of about 5% across the duration of the experiment.

Performance on each measure was also submitted to a split-half analysis to test not just whether performance on each individual measure changes over time, but whether the relative pattern of differences between groups changes also. For this analysis, T2 overall performance (no exclusions) was submitted to a 2x2x2 ANOVA with factors of split-half (first or second), T1-Memory Colour Similarity (Similar or Different) and T2lag (lag2 or lag4).

There were no significant interactions between splitHalf, colourSimilarity and T2Lag ($F(1,80)=0.068$, $p=.795$, $\eta_p^2=.001$), between colourSimilarity and T2Lag ($F(1,80)=0.141$, $p=.708$, $\eta_p^2=.002$) or between colourSimilarity and splitHalf ($F(1,80)=2.996$, $p=.087$, $\eta_p^2=.036$, though note the potentially marginal effect here). There was a significant interaction between lag and splitHalf ($F(1,80)=22.334$, $p<.001$, $\eta_p^2=.218$), which post-hoc comparisons found was driven by performance at Lag2 being significantly worse than at Lag4 in the first half ($p<.001$) but not the second half ($p=.099$). It was also found that performance was overall worse at Lag2 than lag4 ($F(1,80)=8.994$, $p=.004$, $\eta_p^2=.101$) but there was no significant main effect of splitHalf ($F(1,80)=0.146$, $p=.704$, $\eta_p^2=.002$) nor of colourSimilarity ($F(1,80)=3.865$, $p=.053$, $\eta_p^2=.046$, though note the potentially marginal effect here).

However, when the more lenient criteria described above was employed, to account for effects of T1/T2 swapping, all of these differences were weakened, and none reached statistical significance (all $p>.165$)

Discussion

This experiment aimed to investigate whether holding a colour in memory might affect the temporal, rather than spatial, processing of items with the same colour. Overall, no such evidence was found, but there are several issues identified that might preclude a firm conclusion from available data.

RSVP Task

The core analysis, considering T2 response performance on only trials where both the T1 and memory responses were correct, found no differences based on T2's lag nor the similarity between T1's colour and that in memory, and no interaction between the two. What is perhaps most surprising about this is the lack of a difference between the two lags, as one would expect to observe a simple 'blink' effect even if this were unaffected by the contents of memory.

One initial reason for this lack of blink effect could be something as simple as ceiling effects. Performance in all groups was approximately 80%, which is much higher than similar lags in some previous studies (e.g. Chun & Potter, 1995; Raymond et al., 1992). However, other experiments have shown at least weak blink effects with this range of accuracy (Nieuwenstein et al., 2005). Furthermore, even if this were the reason for a lack of differences between conditions, it is unclear what aspect of our design might have led to such ceiling effects.

Although there was no evidence of a 'blink' (poorer T2 performance at lag2 than lag4) in the core analyses, there did appear to be some evidence of one when all trials were included. However, this was specific to the strict T1/T2 criteria and this pattern disappeared when using the liberal criteria. This suggests that even this apparent effect may have been due to a higher chance of confusing the temporal order of T1 and T2, but not necessarily a deficit in the ability to perceive the values correctly. Taken together, these patterns of results suggest a lack of any sort of consistent blink effect, regardless of whether T1 or memory responses were correct or not.

It is not immediately obvious why the current design did not induce a consistent blink. One possibility could be the online nature of the study compared to traditional lab studies, but AB effects in online studies can be comparable to those of lab based studies (Semmelmann & Weigelt, 2017). Key stimuli were all presented centrally, so a lack of control of eye movements should not have had any major impact on the pattern of results, and far from disengaged participants performing poorly, most of our participants seemed to perform well overall. There was a relatively high exclusion rate, but most of the participants excluded did not have consistently poor performance, often performing well in most conditions but perhaps having one that was just below our inclusion threshold. Importantly, when the more liberal criterion was adopted for the RSVP task, the increased numbers of 'valid' trials led to a relatively low exclusion rate. Together, these suggest that participants seemed to be engaged and perform well on what was an admittedly difficult task, with no obvious 'strategy' that might have enabled different behaviour than in traditional lab studies.

Instead, the most obvious design difference to the classic AB design was that the current experiment was a dual task, where something must be maintained in memory while simultaneously performing the RSVP task. When attentional resources are already diverted to a demanding stimulus, such as holding items in working memory, remaining attentional resources may not be sufficient to bias selectively inputs from other secondary stimuli, such that the typical processing advantage for attended stimuli may be weakened or eliminated (Fockert et al., 2001). Holding information in working memory has been shown to weaken or eliminate 'blink' effects relative to when all attention is available to be dedicated to the AB task (Olivers & Nieuwenhuis, 2006). One interpretation from the boost-bounce model of the AB (Olivers & Meeter, 2008) is that with attentional resources committed to maintaining the memory item, remaining resources may not have been sufficient to trigger the overcompensation of selective attention following the presentation of D1. However, other dual-task experiments have shown that it is still possible to observe a 'blink' (e.g. Akyurek & Hommel, 2005; Nieuwenstein et al., 2007).

Another possible reason for the lack of difference in performance between lag conditions may be that the lag positions chosen for analysis were not sufficiently distinct. It is common in previous studies also to sample later lags (such as up to lag8; Raymond et al., 1992), which one could be confident should be after the critical blink period, and therefore act as a baseline against which to detect any blink at earlier lags. This was purposefully not done in the current experiment due to practical limitations on the number of conditions we could include to reach at least a minimum number of desired trials per condition and still keep the overall experiment duration to twenty minutes (given its online nature). To have increased the chances of detecting a blink, the design could have sampled a maximal blink lag (such as lag2) and a minimal blink lag (such as lag7), but the concern was that these extremes might suffer too much from floor or ceiling effects to be sufficiently sensitive to detect any potential influence due to different contents of memory. It was therefore decided to include lag4, which was considered most likely to be during the recovery transition from the blink and be maximally sensitive to memory-guided attentional capture effects. In the current experiment it is possible that lag4 may still have been too early and during a strong blink period, such that sampling a later lag would have revealed more reliable differences. However, if a single later lag is sampled then there is again the risk that it may be too late in the blink recovery to show any strong modulation from memory similarity.

To rule out this possibility in any future studies, it may be important to find ways to sample more lags, allowing one to build up a better picture of the temporal dynamics of the blink for a given sample and design, and increasing the opportunity for detecting any potential differences. A suggestion that may be a compromise between this ideal while still keeping experiment duration to a minimum would be to sample multiple lags but combine them into lag ‘groups’ to ensure sufficient trials per condition for statistical analysis. For example, if one were to sample each lag between 2 and 7, one could then group these into either two groups of three (lags2/3/4 and lags 5/6/7) or three groups of two (lags2/3, lags4/5 and lags6/7). Furthermore, even if each individual lag condition might lack sufficient trials to include each as a separate level in the analysis, even this ‘rough’ data might be useful enough to inform interpretation and establish the temporal dynamics of the blink (or lack of). This approach will be taken in the following experiment in this chapter.

This increase in the spread of sampled lags might also help to protect against any potential issues of temporal predictability. In the current study, T2 was always presented either 200 or 400ms after the onset of T1, and this relative predictability might have allowed temporal attention to operate in a manner that led to such high accuracy and a lack of sensitivity to any further differences from different categories of memory contents. However, there was an additional temporal jitter in the form of T1’s lag appearing at lag 3, 4 or 5. Some evidence in favour of this explanation was the apparent decrease in the relative differences between lags in the 2nd half of the experiment compared to the 1st half of the experiment (when considering all trials, without exclusions). This cannot be explained by simple practice as performance at both lags should improve similarly whereas this effect seems to have been due both to a relative increase in lag2 performance and a decrease in performance in lag4. The latter especially would not be predicted from simple practice effects, and instead might be explained by effects such as temporal predictability, where performance depends less on the processing of T1 itself. The above suggested increase in the range of lags sampled should help to protect against this possibility in any further studies.

However, even this potential ‘blink’ pattern in the first half of trials was no longer evident after accounting for T1/T2 swapping by using the more liberal criteria. This apparent change over time may therefore not be due to any change in sensitivity to T2’s identity but simply an improvement in encoding of the correct temporal order of the two target identities.

Participants appear similarly likely to be able to perceive T2's identity in both halves of the experiment, but with practice they made less errors of which order the two values were presented in.

This issue of T1/T2 swapping appears critical to most of the analyses. Several measures, including T1 performance and T2 performance showed an apparent worse performance when T2 was at lag2 than at lag4, but this difference was generally either weakened or eliminated when using the more liberal criteria. This again implies that for most of these measures, no 'blink' was truly observed as participants seemed to perform equally well at perceiving the correct target identities regardless of whether T2 was at lag2 or lag4, but when T1 and T2 were temporally closer together they were more confusable and there were more trials in which the order they were reported in was incorrect. Note that it is not necessary for both T1 and T2 to be identified correctly for them to 'swap' and it would be similarly possible for the temporal order of one correct target identity and one incorrect target identity (or even a low confidence guess) to be confused when close in time.

If encoding is not strictly sequential, such that both targets end up being encoded with some temporal overlap at short latencies, then their temporal order might be more likely to be confused relative to longer lags where T1 is more likely to be fully encoded before then encoding T2. However, especially given that T1 was always a colour singleton, it remains ambiguous whether these swaps were truly perceptual in nature. Future research might change the question asked from recalling T1 and T2 based on their temporal order, to asking participants to recall the coloured number and the black number (regardless of their order). A further possible factor could be if participants perceived only one of the numbers and were unsure whether it was T1 or T2, they might give the same number for both responses. However, as T1 was always the colour singleton, and participants were informed of this, they should always have been able to tell that if the number they saw was black then it was T2 and otherwise it was T1.

One ambiguous result is the 'marginal' interaction in T1 identification when including all trials which disappeared when using the liberal 'swap' criterion. Marginal results should always be interpreted with caution and would need to be replicated by future research before any firm conclusions could be drawn. However, if replicated in future, this qualitative pattern could suggest that T1/T2 temporal swapping at longer latencies is more likely when T1's colour is different to that in memory. Such a result would in turn suggest that memory contents may indeed affect the processing of presented stimuli but primarily by increasing the temporal resolution of its representation (hence fewer temporal order errors) rather than by improving the speed or accuracy of encoding its identity (as had been predicted). However, the evidence presented here remains insufficient to draw any such conclusion for now.

The only lag effect still to be statistically reliable after taking into consideration temporal swapping was T1 performance when considering only trials on which the memory response was correct (to consider only trials in which there is evidence that the correct colour was being held successfully in memory), though even here the effect was marginal. Although temporal swapping did play some role (as the magnitude and reliability of the apparent difference did decrease when including swapped trials) it cannot fully account for this difference between lag conditions alone.

This remaining difference between lag conditions is fairly surprising. The initial rationale for this study predicted such a lag difference in T2 performance due to different temporal dynamics of a blink triggered by a T1 with either the same or different colour to that being held in memory. Instead T1 performance identification seems to be retroactively affected by a subsequently presented T2. It is not surprising that the perception of a target can be influenced by other subsequent stimulation, as is common in a range of types of masking (e.g. Raab, 1963), however simple low-level perceptual masking effects alone might not account for this effect as the physical characteristics of the subsequent stimulation is following T1 is always a stream of black alphanumeric characters and there is no obvious reason that having a number at T2 should be any more perceptually disruptive at a low level than a letter. Instead, this effect is more likely to be due to higher level interference specific to its nature as a to-be-remembered target rather than a to-be-ignored distractor.

One possible interpretation for this effect comes from a common explanation for the classic blink phenomenon, that T2 cannot be encoded properly if it is presented while encoding resources are still committed to encoding T1. Although the apparent reversal of the effects of competition for encoding resources is unusual considering it is the opposite to previous results (Akyurek & Hommel, 2005; Chun & Potter, 1995; Di Lollo, Kawahara, Shahab Ghorashi, & Enns, 2005; Lunau & Olivers, 2010; Nieuwenhuis & de Kleijn, 2011; Nieuwenstein et al., 2005; Olivers & Meeter, 2008; Raymond et al., 1992; Sergent & Dehaene, 2004; Wyble, Potter, Bowman, & Nieuwenstein, 2011; Zauner et al., 2012), a similar explanation might apply here. It may not be the case that encoding resources always prioritise whichever target was presented first but might instead prioritise either target depending on some factors such as which item has the greater chance of being encoded successfully (and would therefore be a more efficient use of encoding resources). In a typical AB paradigm, the key competition for encoding resources occurs at the presentation of T2. In most studies, at this moment T1 will already have started to undergo encoding mechanisms and might usually be more likely to be prioritised for further encoding, as it is the more likely to be encoded successfully. However, if circumstances mean T1 were comparatively more difficult to encode than a clearer T2 (such as if T1 were less salient) then encoding resources might switch and prioritise encoding T2 instead, to maximise the probability of successfully encoding at least one item. Even if true, it is not obvious which aspect of the current design would account for such a reversal from the traditional direction.

One possibility might be that T1 was always a non-black colour and, though this might be considered to increase its saliency relative to the black T2, it has a luminance closer to the background (low contrast), whereas a black (high contrast) T2 might be more salient, easier to encode, and therefore prioritised. In principle, this direction of encoding competition, though unusual, could still reveal differences due to the influence of T1-Memory colour similarity, such as if T1 having the same colour as that in memory makes T1 easier (or more difficult) to encode, this might still be reflected in differences in performance in identifying T1. This explanation could be explored further by varying characteristics that might affect the relative salience of T1 and T2 and investigate whether this can lead to systematic biasing for one or the other, or whether T1 still takes precedence. However, for the next study it was decided not to try and vary these parameters and instead to keep similar physical characteristics, so that more direct comparisons could be made between the two studies regarding the more likely or more important considerations. The ‘colour singleton’ nature of T1 might be argued to have affected also the ‘blink’. However, other studies have used a salient T1 with different features than the distractors, and still seen a reliable blink at short lags (e.g. (Raymond et al., 1992).

Taken together, these results imply that the temporal dynamics of processing information is unaffected by its similarity to information already in memory. However, the lack of any reliable lag effect makes such a conclusion premature. It remains likely that whatever factors led to the lack of a reliable blink effect might have meant there was no way to detect subtle differences between blink patterns had they been present. It is worth noting though that a reliable AB was not actually a true prerequisite for the focus of this study. It might still have been possible to observe a relative difference in T2’s identification performance based on T1’s similarity to an item held in memory, even in the absence of a reliable ‘blink’. Even if any relative difference in T2’s identification between the two lag positions was not sufficiently reliable (the blink effect), a relative difference in T2's identification based on T1's memory similarity might still have been observed. Therefore, even in the absence of reliable differences between the lag positions, the lack of differences between the memory similarity conditions is still surprising and worthy of further investigation.

One important distinction between this RSVP paradigm and the spatial attention paradigms (e.g. Olivers et al., 2006) is which items are present simultaneously. In most previous versions the key information (the search target and the colour singleton) are presented simultaneously and must compete for attention and focussed processing whereas in the current version only one item is ever visible at once, and always at a central predictable location, requiring no shifting of either covert or overt attention. Therefore, the lack of any memory guided effects here might indicate that memory contents may serve only to bias competition between stimuli but not necessarily affect the processing of a stimulus once attended. However, this conclusion would still not fully account for some previous results showing at least some apparent effects on temporal processing of sequentially presented items from information held in mind (Pashler & Shiu, 1999).

It is also important to consider what information was actually being encoded. In the current study, T1 always had a colour which could be the same or different to the information already encoded to memory, but it was not T1's colour that then needed to be encoded for later recall but its numerical identity. Therefore, even if memory were to bias processing towards stimuli with similar features (e.g. colour), this might not necessarily confer benefits to the processing or encoding of independent feature dimensions (e.g. alphanumeric identity). Previous spatial attention paradigms also use two independent feature dimensions between the feature that identifies a target and the feature that needs to be reported (e.g. the white circle in experiments 6-8 indicating where the target was but the numerical identity needing to be the feature actually reported).

However, in those studies the colour singleton was often irrelevant to the search task and chromatic information itself had no predictive value (e.g. Olivers et al., 2006), whereas in this task, the colour singleton was always a to-be-encoded target and chromatic information could feasibly be used to help alert participants to its onset. In previous spatial tasks, a colour-memory match might help to draw one's attention towards either the correct target (when spatially congruent) or towards a separate distractor (when spatially incongruent), and the additional cost of processing the distractor, disengaging and re-searching for the correct item might have increased the differences between the two conditions. However, in this task, participants were always initially monitoring for the colour singleton item (T1) and even if memory contents did have a slight difference in increasing efficiency of processing towards this item when the same colour than not, this much more subtle difference in required effort between the two conditions would produce far weaker effects and might also partially account for why no such reliable patterns were observed here. It might also still have been expected that when the colour item was the same as that in memory that identification of the numerical value would have been worse if the memory information potentially biased feature-based attention towards colour feature information and away from alphanumeric information, but no evidence of this was found, suggesting instead that memory guided attention might bias primarily spatial attention.

A far more basic consideration is also the nature of the measure. Many of the previous memory guided attention studies measure reaction times as their dependent measure (Hollingworth & Beck, 2016; Olivers, 2008; Olivers et al., 2011; Soto et al., 2006; Zhang et al., 2011, 2010), whereas in this study all targets must be encoded to memory for later report and instead subsequent recall accuracy is measured. It may therefore be that memory guided attention could affect the speed of processing but if parameters used allow even the slower processing condition to reach a similar eventual identification level then any such differences might manifest only in immediate reaction times rather than later accuracy. Future studies might be able to address this by adjusting response requirements such that on any given trial only a single speeded response is required for the RSVP task.

Memory Task

The discussion so far has focussed on the effects of memory contents on the processing of the RSVP task. However, it is also important to consider whether the RSVP task had an influence on the memory task, especially in case any such effects might have had a reciprocal biasing effect through the trial inclusion criteria. For example, the trial would not move on to the memory test stage until both responses were confirmed for the RSVP response stage. This means that the overall memory delay period may vary as a function of the time taken to submit the RSVP responses, perhaps a proxy of uncertainty or difficulty, which may then have a knock-on effect on memory performance. Furthermore, given T1’s colour was identical to the memory item on half of the trials, participants might have been able to attend this strategically to ‘refresh’ their memory representation, which in turn might have interfered with perception of the numerical identity.

Performance on the memory task was better when T1 had a different colour to the memory item compared to when it had the same colour. This is the reverse of what was expected from likely ‘strategic’ explanations and instead indicates that presenting the same colour in the RSVP stream as the memory item impairs rather than improves subsequent recall/identification. This is especially interesting as it is not accompanied by any reliable effects of T1-memory colour similarity on performance on the RSVP task, suggesting this effect is not simply a trade-off between the tasks.

There are examples of similar results in previous dual-task designs (Akyurek & Hommel, 2005; Koelewijn et al., 2008; Nieuwenstein et al., 2007), where identification of a target in the RSVP stream was impaired if it matched information in memory. One explanation for this was because when a given stimulus was associated with the memory task, participants were then slower to form a second association also with the RSVP task when that stimulus was also present during the stream, leading to worse performance. It is also possible that the reverse could occur, such as potentially in the current study, where if on some trials the colour of T1 formed an association with the RSVP task this might have weakened or replaced existing associations with the memory task. Whereas when T1’s colour was different to the memory item, no competition would occur between which task that colour was associated with. This explanation is still problematic for the current experiment though as the colour itself was largely irrelevant to the RSVP task (though it did cue the onset of T1 and therefore also the probable onset of T2 shortly after), so there is even less reason for the colour to form a stronger association with the RSVP task rather than with the memory task (as in the Nieuwenstein et al., 2007 study). Also, in the current study, when the colour was different than the memory colour, it was always highly distinct (it was never a similar shade of the same colour category). Therefore, unlike in the Nieuwenstein et al. (2007) study, where multiple competing letters might become confused in memory with regard to which task they were presented in, in the current experiment the memory response presents options to choose from, all within a single colour category. Therefore, on trials where T1’s colour was different to that in memory, participants would be able to be certain of which colour in their memory was related to the memory task as it would be visible on screen during response, and therefore which other colour in memory must have been related to the RSVP task.

An alternative explanation might be due to colour inhibition during the RSVP task. Colour is mostly irrelevant to the RSVP task and, though it may cue the onset of the targets, any colour would achieve this and there is no need to attend to or encode the specific colour in any given trial. However, given colour’s salience and potentially distracting properties, it is possible participants may have inhibited the colour information to reduce the extent to which it may have drawn attention and encoding resources away from the numerical identity itself. If this were the case though then this inhibition of a given colour might also weaken or impair any existing memory representation of it. Whereas on trials where the colour in the RSVP task is different to the memory item, such an inhibition during the RSVP stream would have had no effect on the subsequent memory task. No 'control' condition was included with no colour singleton, so it is ambiguous whether repeating the same colour impaired memory versus no colour at all, whether presenting a different colour boosted memory compared to no colour at all, or some combination of both. The first seems the more likely as it is unclear why presenting irrelevant (to the memory task) colours which then need to be inhibited else risk competing or interfering existing memory representations, would benefit existing memory representations. However, this would need to be confirmed by future studies by including a no-colour baseline condition.

The above has discussed several potential issues with the current design that would need to be addressed by future research, but it is rarely ideal to make too many changes at once and therefore there needs to be some prioritization of what the immediate focus of a next study should be. The results from the previous experiment remain inconclusive and lacked evidence of any kind of reliable blink at all, meaning no conclusions could reasonably be drawn on how a blink might be affected the contents of WM. It is therefore important to follow this up before moving on to other questions, to consider what might have led to this lack of a blink effect and which changes to prioritise to improve chances of observing a blink in the next study.

Experiment 10: Refined Design for Investigating Effects of Visual Working Memory on Temporal Processing of Sequentially Presented Stimuli

Introduction

The previous section discussed many potential explanations for the unexpected patterns of results, but perhaps the most fundamental finding was that no reliable attentional blink pattern was observed at all (worse T2 identification when at short lags than longer lags). It remains equivocal whether the lack of memory-similarity differences indicates that the contents of memory do not reliably influence the temporal processing of sequentially presented stimuli (at least in relation to the attentional blink) or whether issues with the chosen design and parameters precluded any blink effects and, if addressed, such effects might then show the expected differences in relation to memory contents. Some previous studies have shown that difficult secondary tasks such as a working memory task can ameliorate the attentional blink pattern of effects (Olivers & Nieuwenhuis, 2006), though other studies have found some 'blink' patterns even with a concurrent secondary task (Akyurek & Hommel, 2005; Nieuwenstein et al., 2007). It does seem likely though that in such a dual task design, the blink may be weaker and more sensitive to design choices. The current study therefore aims to build on the previous study and address some of the identified candidate issues, to see whether even a weak but reliable blink pattern can be observed and, if so, whether it shows any differences depending on T1's similarity to memory items.

The first issue that will be addressed is the high levels of trial exclusion and therefore subsequent exclusion of participants. This meant that many participants had fewer trials per condition that had been intended and many participants needed to be excluded. This had the dual consequences of reducing the sensitivity of any analyses to potentially subtle effects and reducing the representativeness of the sample overall. To address this, the total number of trials per participant will be increased by extending the total duration of the session. This was resisted in the previous experiment given the online nature of the study and a desire for such unsupervised participants to maintain engagement throughout. However, by increasing the number of trials, even if exclusion rates due to incorrect responses on the RSVP and/or the memory task increase slightly, this should still leave sufficient trial numbers per condition for more participants, hopefully increasing the sensitivity to any potentially weaker effects.

The other issue this experiment aims to address are the number of lags that will be sampled. In Experiment 9, attempts to keep the total session duration short meant that there were only enough trials to sample two distinct T2 lags. Lag2 and lag4 were chosen as the most likely to demonstrate the predicted effects. One ambiguity with the apparent lack of such effects though was whether these lags were not the ideal choices given the current design and parameters. There were also potential issues with the temporal predictability of using only two lag positions. Increasing the total duration of the experiment allows for more trials to be included. Therefore, the current study will sample 6 different lag positions (lags 2-7) to increase the temporal jitter and unpredictability of each target onset. Although it is not feasible to include sufficient trials per individual lag to include each as a separate level of the factor, these lags can be combined into pairs (early lags 2/3, mid lags 4/5 and late lags 6/7). However, having the data for each individual lag should enable a better overall mapping of the shape of any behavioural effects and a more confident conclusion of whether any 'blink' like effects are observed with our design.

As with experiment 9, it is predicted that overall T2 identification will be worst at early lags, best at late lags, with mid lags somewhere in between the two. It is also predicted that T2 identification will be worse when T1 is of the same colour as the item in memory than when a different colour. Critically, it is predicted that though the feature match between T1 and the memory item should have minimal or no effect in the late lags, performance at the early or especially mid lags will be significantly worse when T1's colour matches that in memory.

Method

Participants

Fifty-seven participants (33 males, 24 females) volunteered via an online recruitment platform (Prolific.co). The intention was to include fifty participants in the final analyses and so any participants that did not meet minimum performance threshold (outlined below) were replaced by additional participants prior to full analyses. The sample was aged between 18 and 68 (M=29.3, SD=8.0) and fifty were right-handed (seven left-handed). Participants were financially reimbursed for their time with £4 (£8 per hour). All procedures were reviewed and approved by the Departmental Ethics Committee.

Materials and Stimuli

Materials and stimuli were identical to Experiment 9, with the following two exceptions. First, the relative lags that T2 could appear at, relative to T1, were expanded from only 2 and 4 (as in experiment 9) to 2, 3, 4, 5, 6 or 7 (all occurring on an equal number of trials). Second, to accommodate these longer lags, the number of items in the RSVP stream was increased from ten to thirteen (ensuring T2 would never be the final character in a stream).

Procedure

The procedure was identical to Experiment 9, with the exception that the total number of trials was increased from 80 to 120.

Design

The design was almost identical to that of Experiment 9, except that the factor of T2 lag was now defined as different lag 'groups', with the three conditions of short-lag (lags 2/3), mid-lags (lags4/5) and long-lag (lags6/7), instead of just the previous conditions of lag2 or lag4.

Data Analysis

As with experiment 9, multiple analyses were run. There were three response types in each trial (T1, T2 and the memory probe), each of which acted as the DV in different analyses. In each analysis, unless otherwise stated, trials were excluded if either of the other responses that trial (other than the DV itself) were incorrect. In all analyses, after any invalid trials were excluded, participants were excluded if they had fewer than 10 trials in any single experimental condition, or an overall average of fewer than 12 trials (note these minimum trial thresholds are higher than Experiment 9).

Each participant's performance was then calculated as the percentage of the remaining valid trials on which the critical DV response was correct. Any participants whose performance in any given experimental condition was more than 2.5 standard deviations from that group mean were excluded. Remaining participants were submitted to a 2x2 repeated measures ANOVA with the factors of lag group (short lags2/3, medium lags4/5 and late lags6/7) and the similarity between T1's colour to that in memory (same colour or different colour).

As with experiment 9, for any analyses of either T1 or T2, two separate response criteria were adopted. The 'strict' criterion was again where a T1 or T2 response was considered correct only if the numerical response matched the target it was given for. However, a second liberal criterion was also adopted such that a T1 or T2 response was considered correct if the corresponding number was given as the response to either target, regardless of which (therefore even if T1's number was given as the T2 response, T1 would still be considered as 'correct' for that trial). This was to allow for the investigation of temporal 'swapping', which seemed prevalent in the previous experiment.

Results

T2 correct performance given T1 correct and memory correct – 3 lag groups.

Like the previous experiment, the primary focus of this study was how T2 performance varied due to its lag position, and due to the similarity between the colour of T1 and the memory item. Therefore, only trials on which correct responses were given to both the T1 and memory item identification were included.

Of the 57 participants who participated, eight were removed because they did not meet the minimum trial number thresholds, and a further five were excluded as outliers, leaving a final N of 44. Table 24 below includes the descriptive statistics for each condition for the remaining participants. Figure 21 below also shows T2 performance at each individual lag (not lag-groups as in the actual analyses) for only those 44 participants included in the main analysis itself.

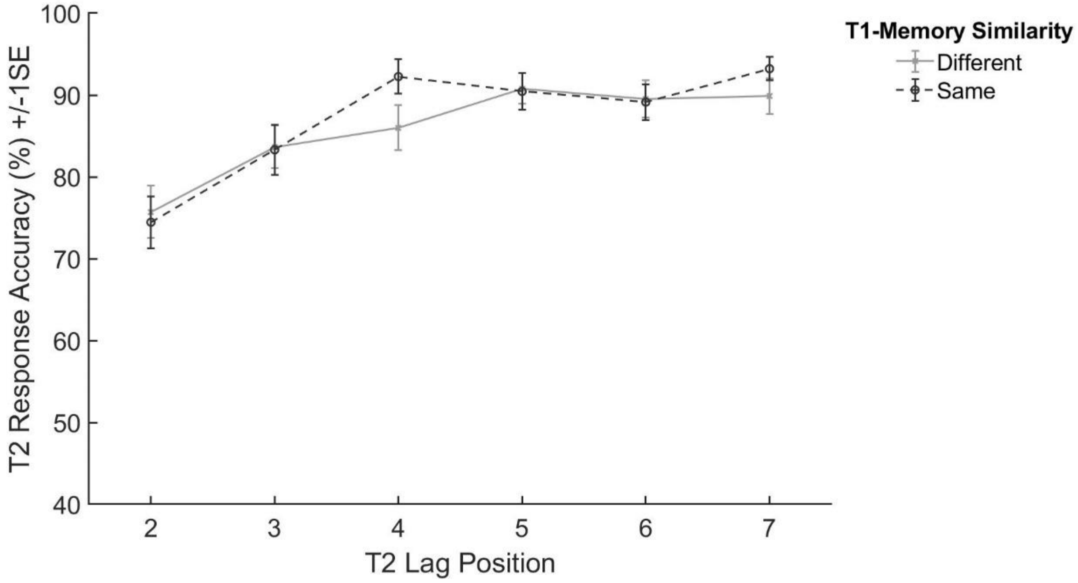

Figure 21: T2 correct performance given T1 correct and memory response correct – individual lag positions.

Table 24: T2 correct performance given T1 correct and memory correct – 3 lag groups. Mean (SD) as %.

		T2 Lag Position			
		Lags2/3	Lags4/5	Lags6/7	Overall
T1-Memory Colour Similarity	Different	80.5 (14.6)	87.9 (13.0)	89.7 (12.7)	86.0 (10.6)
	Same	79.6 (15.7)	91.1 (12.3)	91.1 (9.9)	87.2 (9.0)
	Overall	80.0 (13.8)	89.5 (10.2)	90.4 (8.7)	

Strict Criterion: The overall main effect of lag was statistically reliable (*F*(1.729, 74.355)=23.298, *p*<.001, η_p^2=.351) and pair-wise comparisons confirmed that performance at lags2/3 was significantly worse than lags4/5 (*p*<.001) and lags6/7 (*p*<.001), but that performance was similar at these latter two lag groupings (*p*>.999). Note that the same overall pattern of results is found even if the lags are analysed individually using a 2x6 ANOVA instead, except that lag 3 is not individually significantly different from lags 4 or 7.

However, there was no main effect of T1-Memory colour similarity (*F*(1,43)=1.157, *p*=.288, η_p^2=.026) and no interaction between T1-Memory similarity and T2 lag (*F*(2,86)=0.896,

p=.412, η_p^2=.020). The slight apparent improvement at lag4 when T1's colour was the same as memory is worth mentioning only because it was the pattern predicted since the original experiment. A separate paired-sample t-test found the difference would have been only just significant (p=.041) but this difference is not reliable enough to survive the overall analyses and should not be interpreted as actual positive evidence.

Liberal Criterion: When using the liberal criterion, the additional trials meant an additional six participants were included, giving a total of 50. The main effect of lag was still statistically significant (F(2,98)=23.640, p<.001, η_p^2=.325, again with post-hoc analyses confirming Lags2/3 were significantly worse than Lags4/5 or Lags5/6, both p<.001). This was true if also considering only the same 44 participants in the strict analysis (F(2,86)=21.144, p<.001, η_p^2=.330, with post-hoc comparisons confirming this is driven by Lags2/3 being significantly worse than either Lags3/4 or Lags5/6, both p<.001)

There was also still no reliable main effect of T1-Memory Colour Similarity (F(1,49)=1.759, p=.191, η_p^2=.035, and F(1,43)=0.397,p=.532, η_p^2=.009 for the inclusion of 44 or 50 participants, respectively) and no significant interaction (F(2,98)=0.341, p=.712, η_p^2=.007 and F(2,86)=0.779,p=.462, η_p^2=.0.018, for the inclusion of 44 or 50 participants, respectively).

T1 Correct Given Correct Memory Response

Strict Criteria: T1 accuracy was also analysed as it may also be here that WM may have some effect. Trials were excluded if the memory response was incorrect (though T2's response accuracy was not considered). Following exclusions, 51 participants remained in the analysis, with results shown in Figure 22 and Table 25.

There was a statistically significant main effect of lag (F(2,100)=3.726, p=.028, η_p^2=.069). Bonferroni corrected pairwise comparisons found that this effect was driven by a significantly worse performance at short lags (2/3) than at long lags (6/7; p=.03), but not compared to medium lags (4/5; p=.102). There was also no difference between medium and long lags (p>.999). Furthermore, there was no significant difference between T1's colour similarity to memory (F(1,50)=0.218, p=.643, η_p^2=.004) and no interaction between the two main factors (F(2,100)=0.099, p=.905, η_p^2=.002).

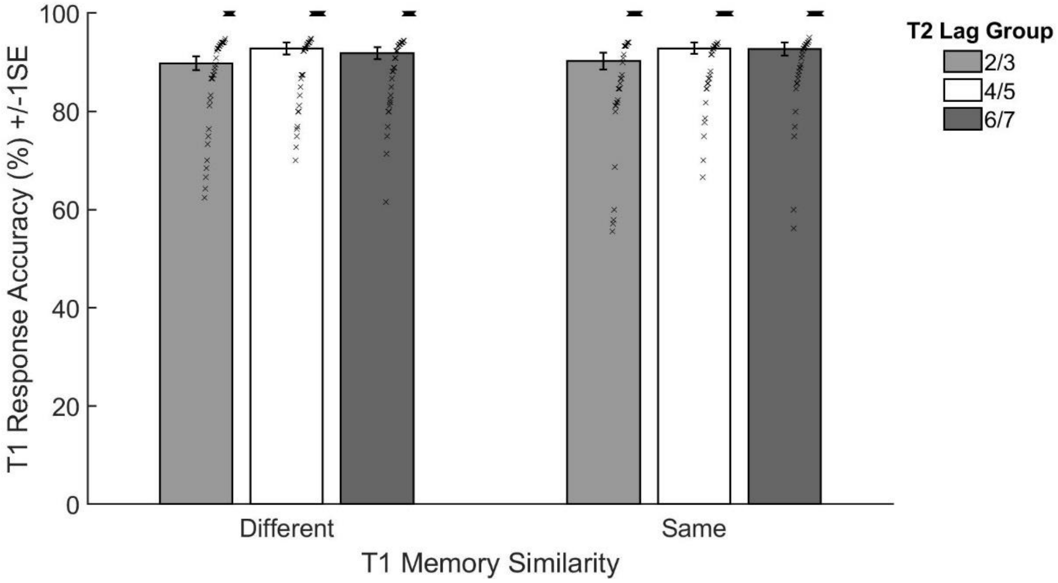

Figure 22: Graph illustrating performance (% correct) at identifying T1, only for trials where the subsequent memory response was correct, split by whether its colour was similar or dissimilar to the memory item, and by T2's lag position. Each x represents an individuals' mean accuracy for that condition (slight horizontal jitter added to display similar results better).

Table 25: Mean (Standard Deviation) performance (% correct) at identifying T1, on trials where memory response was correct, split by T1-memory similarity and by T2's lag.

		T2 Lag Position			
		Lags2/3	Lags4/5	Lags6/7	Overall
T1-Memory Colour Similarity	Different	89.8 (10.4)	92.8 (8.5)	91.9 (8.8)	91.5 (7.6)
	Same	90.2 (12.2)	92.9 (8.6)	92.8 (9.6)	92.0 (7.8)
	Overall	90.0 (9.6)	92.8 (6.9)	92.3 (8.0)	

Liberal Criteria: Despite the more lenient criteria, the number of participants meeting the inclusion was reduced to 49 (because of more participants being classified as outliers) but the pattern is the same even if the same 51 participants from the strict criteria analyses are included. There is then no significant main effect of lag (*F*(2,96)=1.767, *p*=.176**,** η_p^2=.036, or *F*(2,100)=2.046, *p*=.135**,** η_p^2=.039, for the inclusion of 49 or 51 participants, respectively) and still no significant main effect of T1-Memory Colour Similarity (*F*(1,48)=0.306, *p*=.583**,** η_p^2=.006, or *F*(1,50)=0.010, *p*=.921**,** η_p^2<.001, for the inclusion of 49 or 51 participants, respectively) or interaction (*F*(2,96)=0.240, *p*=.787**,** η_p^2=.005 or *F*(2,100)=0.111, *p*=.895**,** η_p^2=.002, for the inclusion of 49 or 51 participants, respectively).

Correct Memory Response Overall

It was also analysed whether T1's similarity to the memory item, or the T1-T2 lag might also cause subsequent differences in performance on the memory task itself. This analysis does not exclude any trials based on T1 or T2 response accuracy, but 2 participants were excluded as outliers, leaving 55 included in the final analysis. Figure 23 and Table 26 below show the descriptive statistics for this analysis.

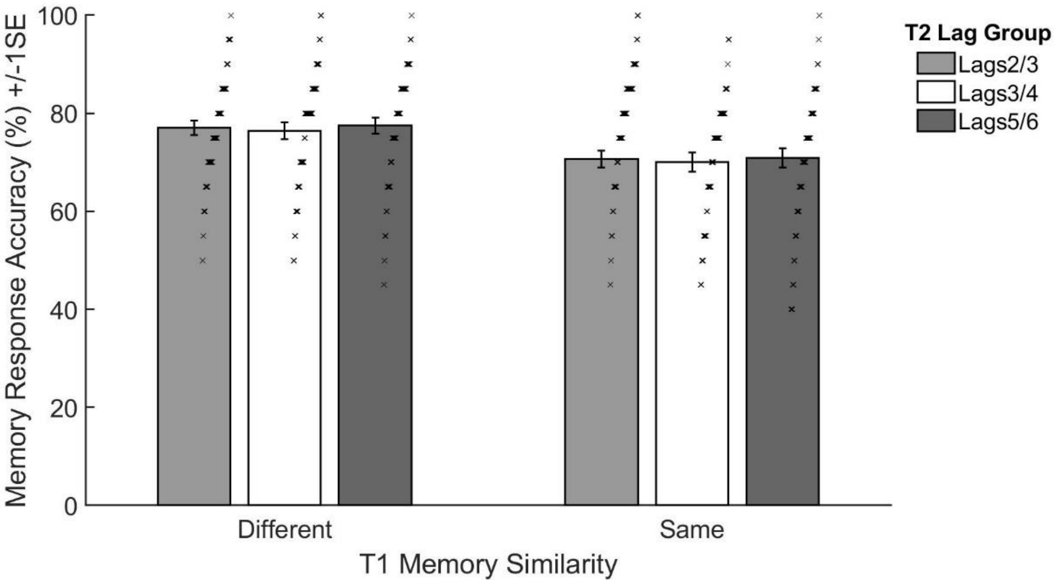

Figure 23: Graph illustrating performance (% correct) on the memory task, split by whether T1's colour was similar or dissimilar to the memory item, and by the temporal lag between T1 and T2.

There was a significant main effect of T1-Memory Colour Similarity (*F*(1,54)=38.265, *p*<.001**,** η_p^2=.415) and no statistically significant main effect of lag group (*F*(2,108)=0.279, *p*=.757**,** η_p^2=.005) and significant interaction (*F*(2,108)=0.003, *p*=.997**,** η_p^2<.001).

Table 26: Mean (Standard Deviation) of performance (% correct) on memory test, split by T1's colour similarity to the memory item, and the T1-T2 temporal lag.

		T2 Lag Position			
		Lags2/3	Lags4/5	Lags6/7	Overall
T1-Memory Colour Similarity	Different	77.0 (10.9)	76.5 (12.8)	77.5 (12.6)	77.0 (9.4)
	Same	70.6 (12.5)	70.0 (14.4)	70.9 (14.5)	70.5 (11.4)
	Overall	73.8 (10.0)	73.2 (11.9)	74.2 (11.5)	

Discussion

This study aimed to build on the results of Experiment 9, to investigate whether addressing some of the identified potential design issues might yield the traditional blink effect and if so, whether its temporal profile varied as a function of the similarity between T1's colour and the memory item. The key design changes were that 6 lags were sampled (though these were combined into pairs for analysis, to improve trials numbers per condition) in attempt to better measure the full temporal profile, rather than sampling only 2 individual lags as in the previous study.

Sampling the additional lags did reveal a potential blink effect such that T2 identification was worse at the shorter lags 2/3 than at the later lags. However, there was still no reliable effect of T1's similarity to the memory item on subsequent identification of T2, either regardless of lag or in any form of interaction. This adds support to the tentative conclusion of Experiment 9, that the contents of memory do not affect the temporal processing or encoding of sequentially presented items in the attentional blink paradigm.

One important difference between these results and those of Experiment 9 are that the apparent differences between the lag positions remained reliable even after accounting for the potential swapping of T1 and T2. Therefore, whereas the initial lag effects in Experiment 9 seemed to be due to an increased likelihood of temporally swapping T1 and T2 when the relative lag was short rather than long, in the current experiment it does seem that identification of T2 was impaired at lags2/3 than at the later lags, more in line with previous AB experiments.

One explanation for why a difference was observed between lags only when a wider range was sampled is due to reduced predictability. In Experiment 9, although the onset of the colour singleton (T1) was slightly jittered, the temporal relationship between T1 and T2 was always either 200 or 400ms and it is possible that this narrow and predictable window of possibilities may have helped to overcome the mechanisms that underly the AB. The less predictable and wider range of T2 lags in the current study though may have been sufficient for temporal anticipation not to counteract the blink mechanisms.

In the previous study, exclusion rates were very high and including only the relatively high performers may have biased the results in a manner that limited the scope for detecting potentially subtle influences from the contents of memory. In the current version, far fewer participants needed to be excluded for poor performance, even despite the slightly higher minimum criterion (the higher minimum threshold was chosen as it had been the originally intended threshold also for the previous experiment, but had needed to be lowered due to the unexpectedly poor performance of so many participants). This sample with participants across a wider range of performance levels may have allowed for sufficient sensitivity to detect the effects from different lags. With this better-quality data and the presence of an attentional blink, these data further suggest that the contents of working memory do not differentially affect the temporal processing of sequentially presented stimuli in relation to whether they share features or not.

Future research should continue to sample wider lags, even if they are then grouped into lag groups for sufficient trial numbers per condition, and careful consideration needs to be given to anticipated trial numbers and performance. Making the dual tasks too challenging risks too many incorrect responses and a high exclusion rate. Making the tasks too easy also risks qualitatively changing the nature of the tasks, such as if the three colour probes were highly distinct colours (such as red, green and blue) and can be remembered or identified based on verbal category labels rather than needing to maintain more subtle visual representations (discussed in more detail in this chapter’s general introduction). Although some research has suggested that memory guided attention effects can spread across broad categories (Moores & Maxwell, 2008; Soto & Humphreys, 2007), other have found that such effects are strongest when requiring participants memory to discriminate within a colour category rather than across categories (Olivers et al., 2006).

Caveats limit the conclusions that can be drawn from this study alone. Lag1 was not sampled. A common result in previous attentional blink paradigms is that when T2 immediately follows

T1, with no distractor stimuli between the two, that T2 identification is relatively unimpaired (e.g. Chun & Potter, 1995). Although referred to as 'lag-1 sparing', it does not appear to be unique to the temporal position itself, but to items in the target category presented without any intervening distractors in a non-target category (e.g. letters or numbers). For example, the blink does not appear to occur if T1 is instead followed by a further two items in the same target category at lag positions 1 and 2 (so no intervening distractor items; (Di Lollo et al., 2005), or if T2 is presented at a similar temporal lag but with no intervening distractor (e.g. just blank; Olivers & Meeter, 2008). Lag1 was not measured as it was planned to group lags together and lag1 would be expected to have too a different level of performance than lag2 or lag3, and so their grouping would have been inappropriate. This lag-1 sparing is not a necessary feature for impaired T2 identification at later lags to indicate a 'blink', but its absence could allow for the alternative explanation by masking, where the physical properties of T1 impair the perception of subsequently presented stimuli regardless of any higher-level information or associations. However, such forward masking effects are unlikely to apply here. T1's chromatic information and therefore brighter luminance gives it a lower contrast with the background, and therefore may be less likely to disrupt the formations of the higher contrast T2 (e.g. di Lollo, 1980). Similarly, adaptation and decreased receptor sensitivity in response to previous stimulation (e.g. Sperling, 1960) would predict that T2s at later lags should show greater impairments than at earlier temporal positions. Even in the absence of measuring T2 identification at lag-1, the current parameters do not appear to make alternative explanations based on masking as more likely than an attentional blink itself. Even if such other factors may play some role, a relative difference may still have been predicted based on T1's memory similarity, and so these would still not account for the lack of observed differences.

Perhaps the main outstanding issue with the current design is that the search targets were not consistent across each trial. In spatial attention paradigms, reliable effects were observed only when the search target was kept consistent across trials, to allow it to be 'offloaded' to long term memory and allow the working memory item to take the role of active search template (Olivers, 2009; Olivers et al., 2011). In the current design, the targets were always 'numbers' rather than a single specific number, which may have precluded memory guided effects. Similarly, in the spatial attention paradigms (Olivers et al., 2006), often the item that participants needed to make a response for in the search task (such an M or an N) was not the stimulus that had the colour singleton property. There was a secondary stimulus that would indicate the location of the search target, such as a diamond (with circles at the non-target locations). This would allow participants to search for this consistent item (e.g. the diamond) even if the target contained might have had multiple potential identities. With a purely sequential search array, an analogous approach could be to have a shape surrounding the RSVP stream which could change properties (such as colour or shape) at the moment of the critical memory-match singleton, to indicate which items in the stream are the targets. This would enable participants to maintain a single constant search template across trials and still allow a wider range of numerical values to be included.

Additional potential improvements would include adding trials where T1 was not a colour singleton, to reduce its relevance to the AB task and preventing participants from relying on the colour singleton to cue the presence of T1. Alternatively, it may be advantageous to vary which item within the RSVP stream is designated the colour singleton.

There is evidence that processing of both T1 and T2 were impaired when the relative lag between them was short. For T1, this seems to be the result of confusing the temporal order of the two targets, as this initial difference was no longer apparent with the liberal criteria. This suggests the processing or perception of T1 itself was not necessarily impaired by a short lag T2, but that even a potentially clearly seen high fidelity representation might have been confused regarding which order the targets appeared. T1 was always the colour singleton though, and T2 was always black, so if the perception of T1 was unaffected then there should be no reason for this confusion (seeing any coloured number should be able to be inferred as being T1). Instead, it is likely that T1's perception was not truly unimpaired but maybe just reached the threshold that participants had an idea of which numbers were presented but not necessarily as full representation with combined featural information, and so unable to use the colour information to distinguish which order the numbers were presented. Even after accounting for the effects of such temporal swapping, the T2 impairment at short lags was still statistically reliable. These 'blink' effect cannot be explained entirely by temporal confusion when the two targets are close together, though this also does not preclude the possibility that at least some such swapping did also occur for T2.

Finally, the improved memory task performance when T1's colour was different to that in memory was replicated again in the current study. The direction of this difference remains surprising as it was predicted memory performance would be better when the memory representation could be 'refreshed' by the colour reappearing during the maintenance interval (in this case as T1). However, as discussed in Experiment 9, there are examples from previous dual-task studies that have found similar effects (Akyurek & Hommel, 2005;

Nieuwenstein et al., 2007). One suggested explanation for such effects is task confusion (Nieuwenstein et al., 2007), such that when the same colour appeared in both the memory and RSVP tasks it might have led to reduced task associations with that colour. Such an explanation still seems unlikely. If T1's colour was present in the memory probe phase then that would always also have been the memory item's colour. Even on trials where participants did not remember the original memory item, their performance would still be improved if they simply chose the response option that matched T1's colour.

A more parsimonious explanation may be that described in the previous discussion, of inhibition. Given that colour is largely irrelevant to the RSVP task, its presence and relative salience might be distracting, especially given that it appears at the time-critical onset of T1. Participants might inhibit this colour information to limit its negative influence on identification of T1, and this inhibition may have a negative impact not only on new incoming information but also on existing internal memory representations. It is possible that all chromatic information may be inhibited to some degree (a future study could test this by using a blocked design and including a condition with no chromatic information in the RSVP stream, to see if this improves subsequent memory performance compared to when there is some chromatic information to be inhibited during the RSVP). But if there is additional inhibition of the actual colour presented, then this would mean that when T1's colour matches that in memory, then the inhibition of this colour specifically may impair the memory representation to a greater extent than when T1's colour is different from that in memory, leading to the observed subsequent difference in performance on the memory task.

General Discussion

These two studies suggest that the contents of memory do not affect either the allocation of temporal attention or the processing of isolated sequentially presented stimuli. This contrasts with previous studies that seemed reliably to demonstrate an influence from the contents of working memory on the allocation of spatial attention. Taken together, the current evidence would suggest that the contents of working memory are able to bias the competition of simultaneously competing inputs towards those that are more similar to featural information already encoded. However, when no such competition is occurring, such as when stimuli are presented in isolation or sequentially as in the current studies, there is no role for any such biasing effect from memory contents.

A key outstanding question would therefore be whether the contents of memory might still play some role in biasing the competition between representations, even when they are not presented simultaneously, under difference circumstances. For example, both the current studies focussed on the similarity between T1 and a memory item, but when T1 is presented, all available processing and encoding resources can be dedicated to it, with no competition from other targets. This might not necessarily be the case with T2 though. If T2 were presented shortly after the offset of T1, and if T1 has not yet been fully encoded, then it may still need to compete with T1 for those limited resources, and maybe here is where the contents of memory could have greater chance of biasing this competition. Future research might therefore wish to adapt the current paradigm but, critically, to vary the feature similarity between T2 and the memory item. This might enable researchers to detect some effects of biasing from the contents of working memory even for sequentially presented stimuli, if it occurs during a critical window of competition for encoding resources.

Many previous AB studies seem to show, or assume, that T1 takes precedent and that if it has not finished being encoded before the onset of T2, that T2 is missed or ‘blinked’. However, experiment 9 presents some evidence that may challenge that assumption. Instead, identification of T1 was impaired when T2 was presented with a short (vs. medium) lag. Note that when T2 had a longer lag, it was still preceded by distractor stimuli, so low-level masking or interference would be similar across conditions, meaning this apparent difference was specific to T2’s identity as a to-be-encoded target. Also, this difference could not be accounted for solely by ‘swapping’ or confusion of the temporal order of the two targets. However, this pattern was not replicated in experiment 2, despite almost identical parameters. To explore this possibility further, future research should try to explore parameters that might mean T1 would not always ‘win’ any such competition. The main theory given here would be that encoding resources should prioritise and emphasise encoding of the most task relevant information possible. In many cases, such as traditional AB studies, this is likely to be T1 as by the time T2 is presented, T1 will already be at least partially encoded and it may be detrimental for encoding mechanisms suddenly to attempt to switch to T2. However, if T1 is a lower salience (such as the lower contrast used in these studies) and is not being encoded reliably by the time a more salient T2 appears, it might be theoretically advantageous for encoding resources then to focus on T2, if it is the more likely to be encoded reliably and correctly, in line with task goals.

Related to this, if the attentional biasing by the contents of memory operates only across simultaneously competing inputs, then the onset of T2 might be the critical moment where such guidance effects might be more likely to be observed. If T2 is presented sufficiently soon after T1 then, even though they are presented sequentially, there may occur competition for encoding resources, and therefore where biasing due to memory similarity may be best observed. It might be thought that T1's similarity would also bias any such competition at the moment of T2's onset in favour of (or against) continued processing of T1. However, the information being encoded in paradigms such as this is the identity, not the memory-related colour itself. Memory similarity may act transiently to draw attention towards a given stimulus. If so, then at the moment of T1's onset, T1 is the only stimulus competing for memory encoding resources, so there is no competition to bias. However, when T2 is presented, its similarity (or dissimilarity) to memory may act to bias encoding resources towards T2, whereas T1's earlier colour may no longer have any such effect. When a 'prime' is presented prior to the RSVP task, T2's performance is worse if it is identical to the prime than different (Koelewijn et al., 2008). Although there was no task demand for this prime to be encoded to memory, these results might nevertheless predict that a similar pattern would be observed if the current paradigm were adapted such that T2 was the critical stimulus. Similarly, the effects did not generally appear to be specific to the 'blink' period, reducing identification of T2 generally, meaning this effect is less likely to be due to changes in temporal processing itself, but perhaps a more generalised effect. In this paradigm the memory item and the search item were identical, with the same information being relevant to both tasks, increasing the possibility of confusion or misattribution (Nieuwenstein et al., 2007). Future research might instead aim to adapt the current paradigm, such that the critical feature of T2 (such as colour) is not the information relevant to the RSVP task itself (the numerical identity).

As has been mentioned previously, if even the traditional 'blink' pattern is not observed, then it becomes more difficult to distinguish whether an associated lack of differences between other conditions (such as memory contents) is truly because no such effects exist, or whether some underlying issue of the parameters or design might be preventing any effects from being observed. This is applicable to experiment 9, which failed to replicate even the expected 'blink' effects and so its lack of effects from different memory match conditions remains ambiguous. However, experiment 10 did replicate a 'blink' effect, with poorer T2 performance at short lags ("200-300ms post T1 onset), and so the continued lack of any memory effects even here presents more compelling evidence that any such memory effects are either absent or too weak to be detected by the current paradigm. Possible differences in the parameters between the two experiments have already been offered that could at least partially account for this difference in results, particularly the first experiment only sampling two lags and having a higher than anticipated rate of exclusion.

An additional avenue for future research might take to investigate whether memory contents may influence temporal processing of sequential stimuli could also be to vary the colour of D1 and investigate whether its similarity to a memory item affects the temporal profile of any subsequent blink. The rationale for this is quite different to that for varying the colours of T1 or T2. As already outlined, the rationale for T1 or T2's similarity to working memory having an impact on the AB was analogous to working memory's effects on spatial attention, biasing the relative degree of attention stimuli receive. However, another way in which WM contents could potentially affect the AB, according the Boost-Bounce model (Olivers & Meeter, 2008), could be if it were to affect distinguishing between targets and distractors. For example, given that the Boost-Bounce theory argues that D1 is critical for the subsequent blink, triggered by its identification as a distractor erroneously benefiting from T1's 'boost', then it is possible that the degree to which D1 is identified as task relevant or a distractor could affect the strength of the subsequent blink. If D1's colour were similar to the contents of WM then it might be classed as less of a distractor and so trigger a reduced magnitude inhibition of a subsequent T2 than when D1's features are dissimilar from those in memory.

If future research builds on the current results by investigating potential interactions between memory contents and either D1 or T2, there are still other outstanding issues that may need to be addressed. For example, there is a potential issue that, although the memory-matching feature (colour) is largely irrelevant to the RSVP task, it could be argued the onset of the coloured item does provide some relevant information as it alerts to the onset of T1. However, that applies for any chromatic information and the specific hue presented is of no further relevance. This was intended to keep the design as similar as possible to the established designs investigating spatial attention. These often also use colour as the memory-matching feature but make the colour irrelevant during the search task itself (usually instead searching for shape information). As described in the discussion of Experiment 10, further efforts could be made to reduce the relevance of this colour singleton to the RSVP task. Key identified methods would be for it to appear also with other items in the stream than T1 on some trials, so that its presence confers no predictive information regarding the RSVP search (such as the onset of one of the targets). Similarly, the temporal position of the targets could be cued by a third feature dimension, as in spatial attention paradigms (Olivers et al., 2006), such as a diamond appearing around T1 and T2 with circles around the distractor items. This would allow participants to adopt a single consistent feature across trials (Olivers, 2009) and would further reduce any potential relevance of the colour itself. A similar consideration regarding the spatial paradigms using a separate stimulus (e.g. surrounding shape) to locate the target than the target itself, is that the target itself does not need to be encoded, merely processed to a sufficient level for a response decision. The guidance of attention from memory therefore seems to exert a greater influence on the guidance of attention to the secondary information (the shape) rather than further effects on the processing of the search target itself. If a similar interplay were to occur in the current design then attention may interact with identifying when a target is present, but may confer no additional benefits to subsequent processing or encoding of the specific identity of the search target itself.

The current design, such as using exact colour matches between the RSVP stream and memory (which might have allowed for colour 'refreshing' strategies) was chosen purposefully to maximise the chances of observing any effects, with the intention of further studies could be utilised to rule out potential confounds. However, it remains possible that this decision may have had the opposite effect and the colour's relevance might have encouraged participants to adopt artificial strategies across the two tasks which may have interfered with the usual mechanisms underlying the AB effect and memory guided attention.

A further consideration though is that AB like effects have been reported previously when using an imagery task rather than a memory task. It is not clear why the imagery version of the task would produce such effects (Pashler & Shiu, 1999) and not the memory version. Similarly, it is surprising that apparent memory effects might occur for loosely associated stimuli such as a word and a line drawing (Sasin & Nieuwenstein, 2016) but not for directly matching stimuli (though as noted previously, this experiment did not include a condition where the line drawing was irrelevant to the memory item). One implication may be that these effects were due to advanced semantic associations and may not occur for simple features such as colour (as used in these experiments). However, this would be markedly different from the apparent parameter requirements of memory driven spatial attention, and potentially may not be due to working memory mechanisms specifically but instead associations in long-term memory.

There are several reasons as to why the spatial memory guided attention effects described in Chapter 4 may not have a temporal equivalent. For example, such studies primarily find a slowing in reaction times, likely due to a delay in attention shifting to the target's location in a serial search, rather than a slowing in the processing of an item once it is being attended. When presented with only one item at a time, every stimulus would always be at the focus of spatial attention. The absence of competition for spatial orienting of attention might therefore preclude observing the memory guided attention phenomenon here. Similarly, as encoding to working memory is argued to be a possible cause of the AB, then if the critical stimulus matched the information already encoded in memory, then it is possible that no further encoding would be required, even if it were treated as a target (rather than a non-target, which would not require any encoding anyway).

Therefore, a major caveat of the conclusions of both Experiments 9 and 10 are that although no evidence was found of the contents of working memory affecting the temporal processing of items, this may only be true of the attentional blink design. As mentioned previously, some studies have found that such a dual task approach can make the attentional blink effect unreliable (Olivers & Nieuwenhuis, 2006). It may therefore be worth exploring alternative measures of temporal attention. One example might be investigating differences in RTs to a target identification task when it is briefly preceded by either a memory-similar or memory-dissimilar stimulus. The preceding stimulus could act either as an enhancing cue or an interfering mask, with the magnitude of either effect mediated by its similarity to the remembered information. This would be similar to the original (Downing, 2000) study, but investigating whether memory-matching stimuli can act as temporal as well as spatial cues.

Unlike in other chapters, the issue of alternative encoding strategies, such as using verbal labels rather than true visual memories when performing the memory task, are less likely to have occurred in Experiments 9 and 10. The RSVP stream explicitly involves alphanumeric characters for which verbalisation would not cause issues, and the memory stimuli use within category discriminations rather than across colour categories, therefore making the application of verbal labels less tenable. However, only five different luminance exemplars were used per colour category. It is therefore possible that over the course of the session, as participants gained increased familiarity with these levels, that they may have been able to apply similar labels (such as a number between 1 and 5). Such non-visual memory representations would therefore not be expected to exhibit the same biasing of visual attention as the expected visual memory representations. The online nature of the study would mean that the common solution of making the stimuli less discriminable may be problematic as it would not be possible to confirm how such smaller differences in intensity would appear on the range of uncalibrated displays used by participants. Similarly, even if a lab-based study were used instead, the inclusion of an articulatory suppression task would also be more problematic than in previous chapters. It might make an already a complex dual task design even more challenging, leading to disengagement and poorer performance, and it would also risk interfering with performance of the RSVP task. Instead, future studies concerned about this possibility might instead introduce jitter to the intensities that could be sampled for each colour. This would allow the continued use of relatively large differences that could be reliably discriminated on users' own displays but still introduce enough variability into the exemplars to prevent participants from becoming overly practiced or familiar, and therefore discouraging such labelling strategies.

As in previous chapters, a final issue is whether the null results reported here could be due to insufficient statistical power. Sample sizes in experiment 9 and 10 were comparable to similar online studies, and greater than many lab-based studies. However, a power analysis would have been a more reliable approach. For the attentional blink effect itself, previous studies have reported consistent effect sizes in the range of η_p^2 between .4 and .6 (Chun & Potter, 1995; Koelewijn et al., 2008; Nieuwenstein et al., 2005; Raymond et al., 1992; Shapiro et al., 2017; including also online based studies such as Semmelmann & Weigelt, 2017). An analysis in G*Power (Erdfelder et al., 2009) revealed that to detect similar effects reliably (with a power of at least 0.8), sample sizes would need to be at least between 9 and 15 participants, so the sample sizes used in this chapter should have been more than enough to detect such effects (50 participants is sufficient for effect sizes down to η_p^2=.15). However, the effect sizes in experiment 9 were much smaller (between η_p^2=.004 and .011), suggesting the null results reflected a genuine weakening or abolishment of the blink effect. Experiment 10, after addressing some of the suggested methodological issues, demonstrated a larger effect (η_p^2=.351) which, though still weaker than previous studies, was found to be statistically reliable.

However, the critical aspect of these studies was not simply replicating the well-established blink effect but determining whether it may be influenced by the contents of working memory. Previous studies reporting effects due to memory or imagery shows that such interactions with the blink effect are weaker (approximately η_p^2=.08, e.g. Nieuwenstein et al., 2007; Sasin & Nieuwenstein, 2016), which would require sample sizes of approximately 100 to detect reliably. The sample sizes used in this chapter (50) may therefore have lacked sufficient sensitivity to detect similarly small effects reliably. However, the observed effect sizes in experiments 9 and 10 were considerably smaller (between η_p^2=.01 and .02). The null results reported cannot therefore be accounted for simply as due to a lack of sufficient power to detect already established effects, and instead suggest a genuine weakening (or absence) of such effects. Overall, mental imagery seems to show a stronger influence on temporal attention than working memory specifically. Future studies would be needed to investigate this further and confirm, and such studies should be mindful of these weaker effects and account for them accordingly.

Chapter 6: General Discussion

Overall Summary

The broad aim of this thesis was to further our understanding of how working memory operates at different levels of representation within a hierarchically structured perceptual system. The series of research focussed on two broad mechanisms: how information calculated across a group of stimuli affects processing of individual items; and how previously encoded information affects the processing of new incoming information. A major focus of the research was how any such mechanisms may operate over time, rather than over static displays of simultaneously presented stimuli, as had previously been used in the majority of the existing research. However, experiments also attempted to adapt existing paradigms for the tactile domain, as well as investigating whether visual feature dimensions that are easier to quantify could be used to investigate working memory processes in more detail. Overall, all experiments failed to find the expected patterns of results. Surprisingly though, almost all also failed to replicate even well-established effects using similar conditions. Overall, the thesis suggests that the previously observed effects may be more limited than previously thought, not occurring reliably when key task conditions are changed. However, it therefore also remains possible that the predicted effects may be observed if using other parameters.

Hierarchical Representations in Working Memory

Chapter 2 aimed to extend on recent research to investigate whether supra-item encoding could be observed also for stimuli presented sequentially, as past literature had primarily focussed on simultaneous presentation. The ultimate intention had been also to distinguish between the integration of spatial information across items presented sequentially, and the actual temporal information itself (such as which order items were presented in, even if at the same single location). Its first attempt was to compare a paradigm similar to previous paradigms demonstrating reliable effects across spatially separated items, when presenting the arrays simultaneously or sequentially. When the non-target items changed values, they changed both their spatial position and temporal order, to maximise the chances of detecting configuration effects due to both/either mechanism, with the intention then to vary changes in spatial and temporal positions separately.

However, Chapter 2 found no evidence of such effects and so the evidence presented here would suggest that the configuration effect (and, by extension, at least some group-level processing) does not occur across stimuli presented sequentially. Importantly, it also failed to replicate previous effects even when using only arrays presented simultaneously in

Experiment 2. As discussed in more detail in Chapter 2, observing effects in the simultaneous condition is not an absolute prerequisite for expecting effects in the sequential condition, but it does raise the possibility that there was some aspect of the task design or parameters that precluded such group level effects, regardless of presentation method, and that if these were addressed that the predicted effects would yet be observed in both simultaneous and sequential presentation conditions. The two key proposed methodological issues of Experiment 1 were the use of a retro-cue (whereas past experiments had primarily used a post-cue), and that array presentation types were mixed and unpredictable within each block of trials (whereas past experiments have often used a single presentation type per trial block, allowing for predictability). However, experiments 2 and 3 found that, even when these issues were addressed, that still no configuration effects were observed. Again, although this therefore suggests that cue type does not interact with the configuration effect, if future research does identify the design or parameter choices in the current experiments that precluded any such configuration effect, such that a reliable effect is observed again, then differences between retro and post cues may yet be revealed. Chapter 3 attempted to extend the configuration effect into the tactile domain, given it was reasoned that tactile perception may have an even greater propensity for integrating information across inputs. Here the focus was on only simple spatial effects across arrays of simultaneously presented stimuli. However, again, no such effects were observed. The main issue identified was that Experiment 4 may simply have been too difficult, but even an easier version in Experiment 5 found no effects.

Overall, the data presented in this thesis suggests that configuration effects over sequentially presented visual stimuli, or simultaneously presented tactile stimuli, is either absent or much weaker than its visuo-spatial counterparts. Furthermore, even the visual spatial effects may be more sensitive to specific stimulus parameters than previously thought. This is surprising though as there is some previous evidence that had suggested such effects should be possible. For example, in the visual domain, some ensemble encoding (a possible candidate mechanism underlying the configuration effect) has been observed across sequentially presented stimuli (Albrecht & Scholl, 2010; Haberman et al., 2009), but such studies often required participants to focus on calculating the summary statistics (such as the 'average') explicitly. The lack of such effects in Chapter 2 while using the configuration effect (that does not require explicit representation of any specific summary statistic) may suggest that these previous results were due to artificial strategies and that such effects may not reliably occur naturally when not the explicit focus of a purposeful strategy. Similarly, the lack of configural effects in Chapter 3 were also surprising given previous studies have demonstrated integration between tactile inputs (Kahrimanovic et al., 2009; Roberts & Humphreys, 2010; Yoshioka et al., 2011), such as the perception of roughness. Experiment 5 confirmed that this lack of effects could not be explained simply by the task being too difficult. Instead, probably reasons for the discrepancy with the current lack of results may be that the parameters adopted here, with palms up and the participant passive in receiving the stimulation, have been suggested to produce weaker integration that other postures and more active exploration (Roberts & Humphreys, 2010). Furthermore, such effects have previously been shown primarily as integration between two inputs, and it may be that when using 4 (Experiment 5) or 6 (Experiment 4) inputs, that even if any such integration does occur, that a subtle change from the cued target changing may not cause a sufficient large change in the summary statistic to be detected reliably.

Taken together, the results of Chapter 2 and Chapter 3 help to reveal the potential limits of the configuration effect and, by extension, potentially the limits of between-item level encoding. The results of the current thesis would suggest that though such effects seem reliable between spatially separate items presented in the visual domain, they may not occur between temporally distinct visual stimuli nor between spatially simultaneous tactile inputs (or at least, if they do occur, this may only be under very restricted conditions). Given that the complexity of between-item information would grow exponentially with the addition of an extra dimension (i.e. time as well as space), it may not be surprising that temporal versions of such mechanisms are not broad or robust as across space, but the results of this thesis suggest that they may be even weaker than predicted. The lack of effects in the tactile domain is more surprising though. The tactile perception system is less complex relative to the visual perceptual system and so it may therefore be thought that ‘complex’ group level calculations will not be performed as reliably. However, these group-level statistics appear to arise earlier in the hierarchy, constituting coarse early mechanisms that can be used for early scene grouping processes (Brady et al., 2017; Hochstein & Ahissar, 2002) and can be calculated even in the absence of representations of individual items. Such between item encoding was therefore expected to be even more prominent in the tactile domain than the visual domain, relative to advanced high-resolution individual representations.

Specific design and parameter considerations are discussed in the respective chapter discussions. However, a broad conclusion that appears to apply to both Chapters is that such ensemble encoding mechanisms are not fully ‘automatic’ and are sensitive to task conditions. Importantly, some group-level encoding has been shown to occur across an entire group but also across sub-groups (Brady & Alvarez, 2011). It is this latter point, that such effects are sensitive to whether items are perceived as part of a 'group' or not, that may partially account for the null results in both Chapters 2 and 3.

One of the previous key examples of temporal information affected working memory performance was that of Olivers & Schreij (2014). Participants were shown an array of coloured circles, but this array would appear on screen by 'sliding' from one of the edges (rather than a more typical sudden onset/offset). Crucially, memory appeared enhanced when the direction of this onset motion was consistent between the encoding memory array and test probe array. This demonstrates evidence of memory for the colour of an item being affected by the consistency of information presented between the encoding and test arrays, as in the configuration effect paradigm, but including temporally distinct information. However, this effect may evidence the integration of the varying states of an individual object across time, rather than of relations across distinct objects (Albrecht et al., 2010). Therefore, it might be that such temporal integration occurs most reliably across stimuli that are perceived to be tightly related to each other, such as being part of the same stimulus, rather than between distinct separate objects presented at different moment in time. A similar consideration may also apply to the tactile domain. In the real world, fingers are unlikely to be presented with separate distinct textures to each separate finger and instead are likely to feel only one or two textures at once. The previous examples of integration across tactile inputs (Kahrimanovic et al., 2009; Roberts & Humphreys, 2010; Van Der Horst et al., 2008) might depend on those inputs being perceived as forming a group for which such summary statistics would therefore be appropriate. The stimuli used in Chapter 3 may have been considered separate stimuli and so any ensemble encoding may have become either eliminated or weakened such that the designs used were not sensitive enough to detect it.

Future research may therefore benefit from testing this explicitly and considering stronger ways for items to be 'grouped'. For information presented across time an example could be a single stimulus that changes its feature over time, rather than separate distinct items. As already mentioned, a similar study has shown temporal integration of a stimulus changing size over time (Albrecht et al., 2010), that experiment required participants to make explicit response about the average size and may have involved artificial strategies. It remains to be established whether similar effects would be observed when a design such as the configuration paradigm is used, where such encoding (if present) is not the explicit focus of the study. Alternatively, for sequential items presented at different spatial locations, grouping could be encouraged if each item formed part of a larger 'whole' stimulus, or presented as a single stimulus moving to different positions, and changing feature values as it reaches/passes each key location. A similar consideration may contribute to the lack of configuration effects observed in the tactile domain in Chapter 3. The textures were similar (different types of textile ribbon) and the consistent onset, offset, and motion, should have helped 'group' them. However, usually fingers would be presented with only one or two textures simultaneously and, even with the use of real-world textures, it remains relatively artificial for each of six fingers to be presented with different textures. It remains possible that tactile integration effects may therefore occur most reliably when distinct inputs can be 'grouped' to form a single overall texture, rather than distinct textures. If this requirement for inputs either spatially across tactile inputs, or temporally across visual inputs, need to be perceived as forming part of a single item, rather than an array of distinct items, this would contrast with the previously observed spatial visual effects which have been observed across arrays of distinct items.

Memory Guided Attention

Chapters 4 and 5 aimed to investigate how existing memory representations may affect the processing of new incoming information. Both chapters adapted a dual-task paradigm that required participants to maintain an item in memory and, during the maintenance period, to perform a secondary unrelated search task. The search task could sometimes include information that matched the memory item, or a different feature value in the same dimension. Even though this information (regardless of similarity to memory) was irrelevant to the search task itself, previous studies have demonstrated that search task performance is disrupted more when memory matching information is presented (Houtkamp & Roelfsema, 2006; Olivers et al., 2006; Soto et al., 2005; Woodman & Luck, 2007).

Previous studies had shown that such effects did not require an exact match between the memory item and the additional information presented during the search task (Olivers et al., 2006), but none had investigated how the magnitude of any such effects varied as a function of the level of similarity between these two critical stimuli. One potential reason is that the effect has been observed primarily using colour as the critical feature dimension (Olivers et al., 2006), which does not lend itself well to quantitative parametrisation of similarity. The focus of Chapter 4 was therefore to investigate whether previously observed memory guided attention effects could be replicated using an alternative feature dimension that could more appropriately measured along a unidimensional continuous scale, specifically angle (either angle of oriented gratings, in Experiments 6 and 8, or of direction of motion in Experiment 7). If such effects could be observed using these feature dimensions, then follow-up studies were planned to investigate how such effects varied at different levels of similarity. Importantly, such evidence might also give some clues as to which levels of representation within the perceptual hierarchy were driving such guidance. For example, within the colour domain, it might be that such effects do not vary much due to absolute similarity, but instead show marked changes when thresholds between categories are crossed (e.g. a similar level of guidance by a shade of red for all other shades of red, but reduced for shades of green), which may imply that such guidance is not driven by the precise representation of that exact shade of red, but a 'coarser' representation of the red category. Similarly, the precision of such effects may be shown to vary with task demands (such as whether the memory task requires a very specific value to be maintained, or a broader category).

However, for both angle of orientation (Experiments 6 and 8) and angle of direction of motion (Experiment 7), no reliable evidence was found for the contents of memory biasing attention. The conclusion of these studies might therefore be that this phenomenon of memory guided attention is a narrow phenomenon observed under only very specific lab conditions and is unlikely to manifest in a meaningful way in the real world. However, the presence of a feature singleton during the search task did not show reliable evidence of attracting attention at all, regardless of its similarity to memory. For example, reaction times were similar regardless of whether the orientation singleton was at the same location as the search target or a distractor location. Furthermore, Experiment 8 confirmed that the null results of Experiment 6 were not simply due to a lack of salience of the critical stimulus during the search task. Therefore, as with Chapters 2 and 3, the possibility remains that these null effects may have been partially due to specific design or parameter choices that, if addressed, would then yield the predicted results.

Given Chapter 4 could not replicate reliable memory guided attention effects using either angle of orientation or direction of motion, it was not possible to use these tools to address the original rationale further, of investigating how such memory guidance effects manifest under more subtle manipulation of the relative differences between critical stimuli, or attempting to investigate which level of hierarchical representation drives such effects. Future research on the topic would therefore benefit from continuing the attempts of this thesis to use non-colour features that can be expressed as continuous measures. It remains possible that the use of such continuous measure features in future research would reveal the originally predicted broader (but potentially more subtle) effects of memory on attentional guidance, such as even when holding multiple items in memory.

Chapter 5 instead returned to the use of colour as the critical feature dimension and a simple 'similar' versus 'dissimilar' categorisation. However, the focus instead was whether the contents of working memory affect only spatial attention (as in previous studies) or also affects the temporal biasing of new stimuli. The dual task paradigm was adapted such that participants performed an attentional blink task (Olivers & Meeter, 2008; Raymond et al., 1992) during the memory maintenance period. This involved searching for and identifying two numbers among letter distractors, all presented one at a time in a sequential stream. The first target (T1) was given an irrelevant colour, and the critical comparison was, when the two targets were close together in time, how identification of each was affected by the relative similarity between the coloured target and the colour in memory. Previous research had shown that holding an image 'in mind' using mental imagery affected the magnitude of the attentional blink (and therefore, affected the temporal processing of new information based on whether it was similar or dissimilar to information already in memory). Given the strong link between visual imagery and visual working memory, it was expected that such effects would be even stronger with memory, where it would be possible to control better the exact values of the critical stimuli. Given even an imagined but never seen image seems able to affect the temporal processing of new information (Pashler & Shiu, 1999), or a line drawing associated with a remembered word (Sasin & Nieuwenstein, 2016), it seems unlikely that such effects are driven by very narrow mechanisms specific only to the exact feature values. Therefore, if such temporal effects were observed then it would open additional avenues to explore the role of hierarchical representations in working memory and attention, with the possibility to ask similar questions as were the rationale for Chapter 4 with regard to what level of representation seems to drive any temporal effects, and how these effects vary as a function of similarity between the critical stimuli.

However, both Experiments 9 and 10 found no reliable evidence of the contents of memory affecting the magnitude of the attentional blink. On the surface these results might therefore imply that the contents of visual memory do not affect temporal processing of subsequent stimuli, and previous tentative evidence using imagery or word meaning instead relies on alternative non-visual mechanisms such as semantic associations. However, Experiment 9 initially found no 'blink' effect at all, such that identification of the second target was similar regardless of whether it was presented with a short or long interval. As with the other chapters, this failure to replicate even a well-established previous effect leaves the above conclusion vulnerable to the criticism that specific design or parameter values may have simply precluded any blink at all and that, if addressed such that a blink were present, that it may then be possible to observe relative differences in its magnitude. Experiment 10 therefore attempted to address the key identified methodological issues of Experiment 9, specifically the limited number of sampled lag positions and the relatively low trial numbers and high resulting participant exclusions. With these changes, a blink was once again observed, with T2 identification worse at shorter than longer lag intervals, but still showed no reliable difference in its magnitude when T1's colour matched that in memory or did not. This final experiment allows for perhaps the most confident interpretation, that even when a blink is observed, that it seems unaffected by T1's similarity to information in memory. This therefore adds more weight to the initial conclusion that the contents of memory affect only spatial, not temporal, attentional processing of subsequently presented stimuli.

However, the item that could share similarity with the memory information was always only T1. One explanation offered for the lack of effects here is that although the contents of memory may positively bias attention in favour of similar information, that behavioural effects would be observed only if this biasing acts during a period where multiple inputs are competing. As explained in Chapter 1, attentional biasing may not always lead to observable behavioural differences if there is no competition to bias. In the current studies, when T1 is presented, there are no other stimuli competing for attentional or encoding resources. At the moment T2 is presented, there may be competition for encoding resources between T1 (especially if at short latencies, such that it may not be fully encoded still) and T2. However, by the time this occurs, the memory relevant information (colour) is no longer visible and may no longer be able to act to bias competition in favour of T1 (T1's colour is irrelevant to the search task response, which requires only the identity to be encoded, so there is little reason to suspect that the colour of T1 would form an integral part of the encoding process). Future research should explore this further by making the second target (T2) the item with the colour singleton rather than the first target (as used in Chapter 5). If this is indeed the moment of maximal competition between targets, then it may have greater potential for demonstrating relative differences between the presence of memory matching information or memory irrelevant information.

A secondary, and unexpected, conclusion from Experiments 9 and 10 was that when the T1-T2 lags were short, that it was T1's identification, not T2's, that was impaired. This is the opposite of most traditionally reported effects, where it is the second target that is inhibited (Olivers & Meeter, 2008; Raymond et al., 1992). The suggested explanation for this is that at the moment T2 is presented, there is competition between the continued ongoing encoding of T1 or a switch to encoding T2 instead. Although previous studies tend to find T1's encoding is given priority, the results of Chapter 5 suggest that this may depend on task conditions and that encoding mechanisms may prioritise the more reliable input. If T1 is less salient (and therefore a weaker signal and less reliable encoding) than T2 (as might be argued to be the case in the current experiments), it may be better for overall task performance to commit resources to a more reliable T2 representation than forego it in favour of a weaker T1 representation. Such an explanation would need to be tested explicitly, by measuring whether processing of T1 relative to T2 varies as a function of their relative salience.

Overall

The overarching theme of the rationale for this thesis was to contribute to other recent efforts to investigate how working memory and attention vary across different levels of representation within the perceptual hierarchy. A particular focus was on whether effects previously observed across space also operate across time, but also whether effects observed in the visual domain can also be observed in the tactile modality, and whether explicit investigation of these different levels can reveal subtle mechanisms missed by traditional working memory paradigms. However, the predicted results across all experiments have been consistently absent.

Especially for Chapters 2 and 4, the original rationale required first establishing the previous effects while using the new adapted paradigms before then going on to use these to ask more advanced and interesting questions. However, the failure to demonstrate even the basic effects meant that the experiments of each chapter were then required to focus on attempting to identify the minor parameter differences that might account for the lack of effects. Though still interesting, the results of these experiments are not able to address any of the original questions of their respective rationales, limiting the scope of any conclusions. In Chapters 2 and 4, a difficulty to replicate basic previous effects also raises important caveats for any conclusions proposed from those previous studies. Minor changes to the design or specific stimulus parameters seem to have precluded any manifestation of each relative effect (configuration effect for Chapter 2, or memory guided attention for Chapters 4), raising important questions about the overall generalisability of the previous research. Chapter 5 was able to address some of the potential methodological issues from Chapter 9 and was able to demonstrate a reliable attentional blink effect in Chapter 10, that still did not differ reliably between whether T1 matched memory contents or not. However, if similar changes to experiments in the other chapters lead to more reliable replication of the previously established effects, then it remains possible that the novel aspects of each design may yet yield the predicted results.

Given the predominantly null findings throughout this thesis, one issue that needs to be addressed is that of statistical power. All the lab-based studies had approximately twenty participants, with the online studies having approximately fifty (with any variation primarily due to differing effects of exclusion criteria). This consistency was chosen as many previous studies involved series of experiments where the number of participants varied wildly with no explanation given. As outlined in each experimental chapter, these participant numbers compared favourably with many previous examples of the respective effects. However, although such participant numbers may have been sufficient for such effects to be observed in general, they may have lacked the power to be confident that such effects would be detected reliably in subsequent experiments such as those reported here.

This is especially true of the novel aspects of each design where the sizes of the effects are less understood. For example, it was surprising that Chapter 2 failed to replicate the configuration effect even when using only simultaneous presentation, despite having similar or increased participant numbers than many previous examples. However, if temporal configuration effects in the sequential condition (if they exist) are weaker than their spatial counterparts, then it may not be reliable to use examples of the latter to estimate the participant numbers required to detect the former reliably. Similarly, it is possible that memory guided attention effects in Chapter 4 may occur even for features such as orientation, but may be weaker than those for colour, and so require greater statistical sensitivity to detect. Therefore, the main conclusion throughout this thesis is not that the predicted effects do not exist, but that some of these previously observed effects may be less robust or generalisable than previously thought, given that even the basic expected effects repeatedly failed to replicate.

It should be noted that the first experiment in each chapter was generally intended as a first step towards more complex designs or more subtle effects. The experiments included throughout this thesis were often designed with conditions to maximally facilitate such effects. For example, in the sequential presentation conditions in Chapter 2, potential follow up studies were planned to investigate the separate effects of changes in spatial positioning and changes in temporal order (regardless of spatial location) independently. These first experiments were designed though to compare conditions of minimal change (identical in both spatial locations and temporal order) and maximal change (different spatial locations and different temporal order). The failure to detect predicted effects even under these conditions may indicate that such effects are much weaker than anticipated and statistical power may need to be a major consideration for future research aiming to investigate similar advanced questions where even more subtle effects seem likely.

With hindsight, it would have been more robust to calculate sample sizes statistically. However, such analyses were run subsequently and are discussed in detail in each respective chapter. Sample sizes were generally large enough to detect effect sizes similar to those reported in comparable previous literature. However, the effect sizes observed in the experiments reported here were all smaller in magnitude. Sample sizes may therefore not have been sufficient to detect the unexpectedly weaker effects observed here, but power alone does not account for this difference from previous reports.

Even in those experiments where sample sizes were towards the lower limit of what would be needed to detect some of the previously reported, effect sizes were again considerably smaller. The effects reported throughout this thesis are therefore not merely similar to those reported in previous studies but with insufficient power leading to false negatives. Instead, such power analyses appear to confirm genuinely weaker (or absent) effects, perhaps due to the proposed methodological considerations. Future research should calculate sample sizes statistically and would benefit from considering the much weaker effect sizes reported throughout this thesis, especially for the novel aspects of each design. They should also address the methodological issues highlighted to strengthen the effects further.

Related to this, it is also worth mentioning the use of online experiments in Chapter 5, which eventually produced the most reliable results in Experiment 10. Although online data collection clearly is not appropriate for all forms of research (in the current thesis, there would obviously have been no way to run the tactile experiments of Chapter 3 online), the evidence in this thesis may further support its use as an important research tool, especially in allowing for greater numbers of participants to be collected quickly. Experiments 9 and 10 were particularly complex dual task designs and, though the predicted effects were not fully observed, overall performance did seem reasonable, demonstrating that even complex task designs can feasibly be run online. However, careful consideration needs to be given to overall duration and trial numbers. Attempts to keep the overall testing session in Experiment 9 led to only a very limited number of trials included and, given the relatively large number of trial exclusions (though comparable to exclusions levels in previously reported lab-based versions) led to many participant exclusions.

One of the key changes for Experiment 10 was to increase the overall duration and number of trials, and this successfully yielded more reliable results. One factor that counted against the use of online studies here was the relatively long trials (each trial included a memory encoding phase, an RSVP presentation phase, then response phases both for the RSVP task and the memory task). However, the data presented here suggests that even with a testing session of 40 minutes, overall performance remained comparable to that of a lab-based version, and there was no observable evidence of performance degrading over time due to fatigue (instead performance appeared to improve over time in Experiment 9).

Similarly, one final consideration is whether participants were allowed sufficient practice at each task to reduce variability and noise in their responses that might otherwise mask true effects. All experiments included a practice block to familiarise participants with the task. These practice blocks were short (Chapter 2: 12 trials; Chapter 3: 3 trials; Chapter 4: 16 trials; Chapter 5: 5 trials) but practice blocks were also short in previous studies reporting positive effects for configuration effects (20 trials; Vidal et al., 2005), memory guided attention (16 trials; Olivers et al., 2006) and attentional blink effects (20 trials; Chun & Potter, 1995). In all experiments, the experimenter monitored the practice blocks to confirm participants had understood the task instructions and added clarification or answered questions where necessary, and all participants confirmed that they understood the task prior to commencing the experimental blocks. It is unlikely that the null results reported in Chapters 2 (configuration effect; 12 trial practice block) or Chapter 4 (memory guided attention) could easily be explained by lacking a few extra practice trials. However, the practice sessions in both Chapter 3 (tactile configuration effect) and Chapter 5 (memory guided temporal attention) were admittedly much shorter than might have been ideal. This was due to the pressure to keep the overall testing sessions short, while maximising the number of trials included in the analyses. On balance, given that overall task performance was reasonable, and the dependent measures were mainly accuracy, it seems unlikely that a few additional practice trials would have altered the pattern of effects drastically and instead, as in the transition from experiment 9 to experiment 10, longer testing sessions and instead more trials included in the final analyses may have been of greater benefit. However, some of the more subtle or novel predicted effects may not naturally exhibit themselves sufficiently reliably and may require extended training sessions to be detected. It is also worth acknowledging that experiments in Chapter 4 (memory guided attention) used reaction times rather than simple accuracy, which are more sensitive to noise and variability in responses, and so may have benefited from an extended training session. This is especially true as many participants were from other disciplines or the public, so may not have been sufficiently practiced with cognitive experiments in general.

Ultimately, the evidence presented in this thesis remains ambiguous. The key contribution has been to further the understanding of some of the limits of how far previously observed effects may generalise to. However, there remain too many caveats or ambiguities in the results to draw definite firm conclusions from this data alone, and future research will be needed to confirm some of the proposed explanations, and to identify the critical parameter or design differences to account for the differences in results with those of past studies.

References

Aagten-Murphy, D., & Bays, P. M. (2019). Independent working memory resources for egocentric and allocentric spatial information. *PLoS Computational Biology*, *15*(2), 1–20. https://doi.org/10.1371/journal.pcbi.1006563

Ahmad, J., Swan, G., Bowman, H., Wyble, B., Nobre, A. C., Shapiro, K. L., & McNab, F. (2017). Competitive interactions affect working memory performance for both simultaneous and sequential stimulus presentation. *Scientific Reports*, *7*(1), 1–16. https://doi.org/10.1038/s41598-017-05011-x

Akyurek, E. G., & Hommel, B. (2005). Short-term memory and the attentional blink: Capacity versus content. *Memory & Cognition*, *33*(4), 654–663. https://doi.org/10.3758/BF03195332

Albrecht, A. R., Scholl, B. J., Albrecht, A. R., & Scholl, B. J. (2010). Perceptually Averaging in a Continuous Visual World: Extracting Statistical Summary Representations Over Time. *Psychological Science*, *21*(4), 560–567.

Alilović, J., Timmermans, B., Reteig, L. C., Van Gaal, S., & Slagter, H. A. (2019). No Evidence that Predictions and Attention Modulate the First Feedforward Sweep of Cortical Information Processing. *Cerebral Cortex*, *29*(5), 2261–2278. https://doi.org/10.1093/cercor/bhz038

Alvarez, G. A., & Oliva, A. (2008). The representation of simple ensemble visual features outside the focus of attention. *Psychological Science*, *19*(4), 392–398. https://doi.org/10.1167/7.9.129

Alvarez, G. A., & Oliva, A. (2009). Spatial ensemble statistics are efficient codes that can be represented with reduced attention.

Anwyl-Irvine, A., Massonnié, J., Flitton, A., Kirkham, N., & Evershed, J. (2018). Gorilla in our MIDST: An online behavioral experiment builder. *BioRxiv*, (April 2019), 388–407. https://doi.org/10.1101/438242

Ariely, D. (2001). Seeing Sets: Representation by Statistical Properties. *Psychological Science*, *12*(2), 157–162. https://doi.org/10.1111/1467-9280.00327

Astle, D. E., Nobre, A. C., & Scerif, G. (2010). Subliminally presented and stored objects capture spatial attention. *Journal of Neuroscience*, *30*(10), 3567–3571. https://doi.org/10.1523/JNEUROSCI.5701-09.2010

Auvray, M., Gallace, A., Hartcher-O'Brien, J., Tan, H. Z., & Spence, C. (2008). Tactile and visual distractors induce change blindness for tactile stimuli presented on the fingertips. *Brain Research*, *1213*, 111–119. https://doi.org/10.1016/j.brainres.2008.03.015

Auvray, M., Gallace, A., & Spence, C. (2011). Tactile short-term memory for stimuli presented on the fingertips and across the rest of the body surface. *Attention, Perception, and Psychophysics*, *73*(4), 1227–1241. https://doi.org/10.3758/s13414-011-0098-6

Awh, E., Barton, B., & Vogel, E. K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science*, *18*(7), 622–628. https://doi.org/10.1111/j.1467-9280.2007.01949.x

Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*(8), 437–443. https://doi.org/10.1016/j.tics.2012.06.010

Baars, B. J., & Franklin, S. (2003). How conscious experience and working memory interact. *Trends in Cognitive Sciences*, *7*(4), 166–172. https://doi.org/10.1016/S1364-6613(03)00056-1

Bach-y-Rita, P. (2004). Tactile sensory substitution studies. *Annals of the New York Academy of Sciences*, *1013*, 83–91. https://doi.org/10.1196/annals.1305.006

Baddeley, A. D., & Hitch, G. J. (1974). Working Memory. In *Psychology of Learning and Motivation* (pp. 47–89). https://doi.org/10.1016/S0079-7421(08)60452-1

Bahle, B., Thayer, D. D., Mordkoff, J. T., & Hollingworth, A. (2019). The Architecture of Working Memory: Features From Multiple Remembered Objects Produce Parallel, Coactive Guidance of Attention in Visual Search. *Journal of Experimental Psychology: General*, *149*(5), 967–983. https://doi.org/10.1037/xge0000694

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, Paul M, & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science (New York, N.Y.)*, *321*(5890), 851–854. https://doi.org/10.1126/science.1158023

Beck, Valerie, M., Hollingworth, A., & Luck, Steven, J. (2012). Simultaneous Control of Attention by Multiple Working Memory Representations. *Psychological Science*, *23*(8), 887–898.

Blalock, L. D., & Clegg, B. A. (2010). Encoding and representation of simultaneous and sequential arrays in visuospatial working memory. *Quarterly Journal of Experimental Psychology (2006)*, *63*(5), 856–862. https://doi.org/10.1080/17470211003690680

Bliss, I., & Hämäläinen, H. (2005). Different working memory capacity in normal young adults for visual and tactile letter recognition task. *Scandinavian Journal of Psychology*, *46*(3), 247–251. https://doi.org/10.1111/j.1467-9450.2005.00454.x

Block, N. (2011). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, *15*(12), 567–575. https://doi.org/10.1016/j.tics.2011.11.001

Boduroglu, A., & Shah, P. (2006). Configural representations in spatial working memory. *Cognitive Processing*, *7*(SUPPL. 1), 102–124. https://doi.org/10.1007/s10339-006-0120-4

Boduroglu, A., & Shah, P. (2009). Effects of spatial configurations on visual change detection: An account of bias changes. *Memory and Cognition*, *37*(8), 1120–1131. https://doi.org/10.3758/MC.37.8.1120

Boehler, C. N., Schoenfeld, M. A., Heinze, H. J., & Hopf, J. M. (2008). Rapid recurrent processing gates awareness in primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(25), 8742–8747. https://doi.org/10.1073/pnas.0801999105

Bor, D., & Seth, A. K. (2012). Consciousness and the prefrontal parietal network: Insights from attention, working memory, and chunking. *Frontiers in Psychology*, *3*(MAR), 1–14. https://doi.org/10.3389/fpsyg.2012.00063

Brady, T. F., & Alvarez, G. A. (2011). Hierarchical Encoding in Visual Working Memory : Ensemble Statistics Bias Memory for Individual Items.

https://doi.org/10.1177/0956797610397956

Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity : Beyond individual items and toward structured representations, *11*, 1–34. https://doi.org/10.1167/11.5.4.Introduction

Brady, T. F., Konkle, T., & Alvarez, G. A. (2013). Real-World Objects Are Not Represented as Bound Units : Independent Forgetting of Different Object Details From Visual Memory, *142*(3), 791–808. https://doi.org/10.1037/a0029649

Brady, T. F., Shafer-skelton, A., Alvarez, G. A., Brady, T. F., Shafer-skelton, A., & Alvarez, G. A. (2017). Global Ensemble Texture Representations are Critical to Rapid Scene Perception.

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. https://doi.org/10.1073/pnas.1520027113

Brady, T. F., & Tenenbaum, J. B. (2013). A probabilistic model of visual working memory: Incorporating higher order regularities into working memory capacity estimates. *Psychological Review*, *120*(1), 85–109. https://doi.org/10.1037/a0030779

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436. https://doi.org/10.1163/156856897X00357

Brand, A., & Bradley, M. T. (2012). Assessing the Effects of Technical Variance on the Statistical Outcomes of Web Experiments Measuring Response Times. *Social Science Computer Review*, *30*(3), 350–357. https://doi.org/10.1177/0894439311415604

Busch, N. A., & VanRullen, R. (2010). Spontaneous EEG oscillations reveal periodic sampling of visual attention. *Proceedings of the National Academy of Sciences*, *107*(37), 16048–16053. https://doi.org/10.1073/pnas.1004801107

Chang, D., Nesbitt, K. V., & Wilkins, K. (2007). The gestalt principles of similarity and proximity apply to both the haptic and visual grouping of elements. *Conferences in Research and Practice in Information Technology Series*, *64*(January 2007), 79–86.

Chong, S. C., & Treisman, A. (2003). Representation of statistical properties. *Vision Research*, *43*(4), 393–404. https://doi.org/10.1016/S0042-6989(02)00596-5

Chong, S. C., & Treisman, A. (2005). Statistical processing: Computing the average size in perceptual groups. *Vision Research*, *45*(7), 891–900. https://doi.org/10.1016/j.visres.2004.10.004

Chun, M. M. (2011). Visual working memory as visual attention sustained internally over time. *Neuropsychologia*, *49*(6), 1407–1409. https://doi.org/10.1016/j.neuropsychologia.2011.01.029

Chun, M. M., & Potter, M. C. (1995). A Two-Stage Model for Multiple Target Detection in Rapid Serial Visual Presentation. *Journal of Experimental Psychology: Human Perception and Performance*, *21*(1), 109–127. https://doi.org/10.1037/0096-1523.21.1.109

Clevenger, P. E., & Hummel, J. E. (2014). Working memory for relations among objects. *Attention, Perception, & Psychophysics*, *76*(December 2013), 1933–1953. https://doi.org/10.3758/s13414-013-0601-3

Cole, G. G., Kentridge, R. W., & Heywood, C. A. (2004). Visual salience in the change detection paradigm: The special role of object onset. *Journal of Experimental Psychology: Human Perception and Performance*, *30*(3), 464–477. https://doi.org/10.1037/0096-1523.30.3.464

Connor, C. E., & Johnson, K. O. (1992). Neural coding of tactile texture: comparison of spatial and temporal mechanisms for roughness perception. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *12*(9), 3414–3426.

Corbett, J. E. (2017). The Whole Warps the Sum of Its Parts: Gestalt-Defined-Group Mean Size Biases Memory for Individual Objects. *Psychological Science*, *28*(1), 12–22. https://doi.org/10.1177/0956797616671524

Corbett, J. E., & Melcher, D. (2014). Characterizing ensemble statistics: mean size is represented across multiple frames of reference. *Attention, Perception & Psychophysics*, *76*(3), 746–758. https://doi.org/10.3758/s13414-013-0595-x

Corbett, J. E., & Oriet, C. (2011). The whole is indeed more than the sum of its parts: Perceptual averaging in the absence of individual item representation. *Acta Psychologica*, *138*(2), 289–301. https://doi.org/10.1016/j.actpsy.2011.08.002

Coull, J. T., & Nobre, A. C. (1998). Where and when to pay attention: The neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. *Journal of Neuroscience*, *18*(18), 7426–7435. https://doi.org/10.1523/jneurosci.18-18-07426.1998

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

Cowan, N., Elliott, E. M., Saults, S. J., Morey, C. C., Mattox, S., Hismjatullina, A., & Conway, A. R. A. (2005). On the capacity of attention: Its estimation and its role in working memory and cognitive aptitudes. *Cognitive Psychology*, *51*(1), 42–100. https://doi.org/10.1016/j.cogpsych.2004.12.001

Cowan, N., Saults, J. S., & Clark, K. M. (2015). Exploring age differences in visual working memory capacity: Is there a contribution of memory for configuration? *Journal of Experimental Child Psychology*, *135*, 72–85. https://doi.org/10.1016/j.jecp.2015.03.002

Dakin, S. (2014). Seeing Statistical Regularities. *Oxford Handbook of Perceptual Organization*, (November), 1–15. https://doi.org/10.1093/oxfordhb/9780199686858.013.054

de Gardelle, V., Sackur, J., & Kouider, S. (2009). Perceptual illusions in brief visual presentations. *Consciousness and Cognition*, *18*(3), 569–577. https://doi.org/10.1016/j.concog.2009.03.002

De Gardelle, V., & Summerfield, C. (2011). Robust averaging during perceptual judgment. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(32), 13341–13346. https://doi.org/10.1073/pnas.1104517108

de Leeuw, J. R., & Motz, B. A. (2016). Psychophysics in a Web browser? Comparing response times collected with JavaScript and Psychophysics Toolbox in a visual search task. *Behavior Research Methods*, *48*(1), 1–12. https://doi.org/10.3758/s13428-015-0567-2

Delvenne, J.-F. ßis, & Bruyer, R. (2006). A configural effect in visual short-term memory for features from different parts of an object. *The Quarterly Journal of Experimental Psychology*, *59*(9), 1567–1580. https://doi.org/10.1080/17470210500256763

Desimone, R., & Duncan, J. (1995). Neural Mechanisms of Selective Visual. *Annual Review of Neuroscience*, *18*(1), 193–222. https://doi.org/10.1146/annurev.ne.18.030195.001205

di Lollo, V. (1980). Temporal integration in visual memory. *Journal of Experimental Psychology: General*, *109*(1), 75–97. https://doi.org/10.1037/0096-3445.109.1.75

Di Lollo, V., Kawahara, J. I., Shahab Ghorashi, S. M., & Enns, J. T. (2005). The attentional blink: Resource depletion or temporary loss of control? *Psychological Research*, *69*(3), 191–200. https://doi.org/10.1007/s00426-004-0173-x

Downing, P.E. (2000). Interactions between visual working memory and visual attention. *Psychological Science*.

Downing, Paul E., & Dodds, C. M. (2004). Competition in visual working memory for control of search. *Visual Cognition*, *11*(6), 689–703. https://doi.org/10.1080/13506280344000446

Duncan, J., Parr, A., Woolgar, A., Thompson, R., Bright, P., Cox, S., … Nimmo-Smith, I. (2008). Goal Neglect and Spearman's g: Competing Parts of a Complex Task. *Journal of Experimental Psychology: General*, *137*(1), 131–148. https://doi.org/10.1037/0096-3445.137.1.131

Erdfelder, E., Faul, F., Buchner, A., & Lang, A. G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, *41*(4), 1149–1160. https://doi.org/10.3758/BRM.41.4.1149

Eriksen, C., & St James, J. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics*, *40*(4), 225–240.

Evans, P. M., Craig, J. C., & Rinker, M. A. (1992). Perceptual processing of adjacent and nonadjacent tactile nontargets. *Perception & Psychophysics*, *52*(5), 571–581. https://doi.org/10.3758/BF03206719

Fockert, J. W. De, Rees, G., Frith, C. D., Lavie, N., Science, S., Series, N., & Mar, N. (2001). The Role of Working Memory in Visual Selective Attention Published by : American Association for the Advancement of Science Stable URL : http://www.jstor.com/stable/3082574 Linked references are available on JSTOR for this article : You may need to log in, *291*(5509), 1803–1806.

Fougnie, D., & Alvarez, G. A. (2011). Object features fail independently in visual working memory : Evidence for a probabilistic feature-store model, *11*, 1–12. https://doi.org/10.1167/11.12.3.Introduction

Fougnie, D., Kanabar, A., Brady, T., & Alvarez, G. (2016). Asymmetric confidence intervals reveal hidden information in working memory. *Journal of Vision*, *16*(12), 34. https://doi.org/10.1167/16.12.34

Fu, S., Fedota, J., Greenwood, P., & Parasuraman, R. (2010). Dissociation of visual C1 and P1 components as a function of attentional load: An event-related potential study. *Biological Psychology*, *85*(1), 171–178. https://doi.org/10.1016/j.biopsycho.2010.06.008

Gallace, A., Auvray, M., Tan, H. Z., & Spence, C. (2006). When visual transients impair tactile change detection: A novel case of crossmodal change blindness? *Neuroscience Letters*, *398*(3), 280–285. https://doi.org/10.1016/j.neulet.2006.01.009

Gallace, A., & Spence, C. (2011). To what extent do gestalt grouping principles influence tactile perception? *Psychological Bulletin*, *137*(4), 538–561. https://doi.org/10.1037/a0022335

Gallace, A., Tan, H. Z., Haggard, P., & Spence, C. (2008). Short term memory for tactile stimuli. *Brain Research*, *1190*(1), 132–142. https://doi.org/10.1016/j.brainres.2007.11.014

Gallace, A., Tan, H. Z., & Spence, C. (2006). The failure to detect tactile change: A tactile analogue of visual change blindness. *Psychonomic Bulletin and Review*, *13*(2), 300–303. https://doi.org/10.3758/BF03193847

Gallace, A., Tan, H. Z., & Spence, C. (2008). Can tactile stimuli be subitised? An unresolved controversy within the literature on numerosity judgments. *Perception*, *37*(5), 782–800. https://doi.org/10.1068/p5767

Gallace, A., Zeeden, S., Röder, B., & Spence, C. (2010). Lost in the move? Secondary task performance impairs tactile change detection on the body. *Consciousness and Cognition*, *19*(1), 215–229. https://doi.org/10.1016/j.concog.2009.07.003

Gao, Z., Ding, X., Yang, T., Liang, J., & Shui, R. (2013). Coarse-to-Fine Construction for High-Resolution Representation in Visual Working Memory. *PLoS ONE*, *8*(2). https://doi.org/10.1371/journal.pone.0057913

Gao, Z., Gao, Q., Tang, N., Shui, R., & Shen, M. (2015). Organization principles in visual working memory: Evidence from sequential stimulus display. *Cognition*, *146*, 277–288. https://doi.org/10.1016/j.cognition.2015.10.005

Gilbert, C. D., & Sigman, M. (2007). Brain States: Top-Down Influences in Sensory Processing. *Neuron*, *54*(5), 677–696. https://doi.org/10.1016/j.neuron.2007.05.019

Gilchrist, A. L., & Cowan, N. (2014). A two-stage search of visual working memory: investigating speed in the change-detection paradigm. *Attention, Perception, & Psychophysics*, *76*(7), 2031–2050. https://doi.org/10.3758/s13414-014-0704-5

Gmeindl, L., Nelson, J. K., Wiggin, T., & Reuter-Lorenz, P. a. (2011). Configural representations in spatial working memory: modulation by perceptual segregation and voluntary attention. *Attention, Perception & Psychophysics*, *73*(7), 2130–2142. https://doi.org/10.3758/s13414-011-0180-0

Gorea, A., Belkoura, S., & Solomon, J. A. (2014). Summary statistics for size over space and time. *Journal of Vision*, *14*(9), 22–22. https://doi.org/10.1167/14.9.22

Griffin, I. C., Miniussi, C., & Nobre, A. C. (2001). Orienting attention in time. *Frontiers in Bioscience : A Journal and Virtual Library*, *6*(12), D660–D671. https://doi.org/10.2741/Griffin

Gronau, N., & Shachar, M. (2014). Contextual integration of visual objects necessitates attention. *Attention, Perception & Psychophysics*, *76*(3), 695–714. https://doi.org/10.3758/s13414-013-0617-8

Haberman, J., Brady, T. F., & Alvarez, G. A. (2015). Individual Differences in Ensemble Perception Reveal Multiple , Independent Levels of Ensemble Representation, *144*(2), 432–446.

Haberman, J., Harp, T., & Whitney, D. (2009). Averaging facial expression over time. *Journal of Vision*, *9*(11), 1.1-13. https://doi.org/10.1167/9.11.1

Haberman, J., & Whitney, D. (2011). Efficient summary statistical representation when change localization fails. *Psychonomic Bulletin and Review*, *18*(5), 855–859.

https://doi.org/10.3758/s13423-011-0125-6

Haque, H., Lobier, M., Palva, J. M., & Palva, S. (2020). Neuronal correlates of full and partial visual conscious perception. *Consciousness and Cognition*, *78*(December 2019), 102863. https://doi.org/10.1016/j.concog.2019.102863

Harris, J. A., Harris, I. M., & Diamond, M. E. (2001). The topography of tactile working memory. *Journal of Neuroscience*, *21*(20), 8262–8269. https://doi.org/10.1523/jneurosci.21-20-08262.2001

Hilbig, B. E. (2016). Reaction time effects in lab- versus Web-based research: Experimental evidence. *Behavior Research Methods*, *48*(4), 1718–1724. https://doi.org/10.3758/s13428-015-0678-9

Hillyard, S. A., & Anllo-Vento, L. (1998). Event-related brain potentials in the study of visual selective attention. *Proceedings of the National Academy of Sciences of the United States of America*, *95*(3), 781–787. https://doi.org/10.1073/pnas.95.3.781

Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, *36*(5), 791–804. https://doi.org/10.1016/S0896-6273(02)01091-7

Hollingworth, A. (2006). Visual memory for natural scenes: Evidence from change detection and visual search. *Visual Cognition*, *14*(4–8), 781–807. https://doi.org/10.1080/13506280500193818

Hollingworth, A., & Beck, V. M. (2016). Memory-Based Attention Capture when Multiple Items are Maintained in Visual Working Memory. *Journal of Experimental Psychology. Human Perception and Performance*, *42*(7), 319–335. https://doi.org/10.1037/xhp0000230

Houtkamp, R., & Roelfsema, P. R. (2006). The effect of items in working memory on the deployment of attention and the eyes during visual search. *Journal of Experimental Psychology. Human Perception and Performance*, *32*(2), 423–442. https://doi.org/10.1037/0096-1523.32.2.423

Howard, C. J., Arnold, C. P. A., & Belmonte, M. K. (2017). Slower resting alpha frequency is associated with superior localisation of moving targets. *Brain and Cognition*, *117*, 97–107. https://doi.org/10.1016/J.BANDC.2017.06.008

Hubert-wallander, B., & Boynton, G. M. (2015). Not all summary statistics are made equal : Evidence from extracting summaries across time. *Journal of Vision*, *15*(2015), 1–12. https://doi.org/10.1167/15.4.5.doi

Ihssen, N., Linden, D. E. J., & Shapiro, K. L. (2010). Improving visual short-term memory by sequencing the stimulus array. *Psychonomic Bulletin and Review*, *17*(5), 680–686. https://doi.org/10.3758/PBR.17.5.680

Indovina, I., & MacAluso, E. (2007). Dissociation of stimulus relevance and saliency factors during shifts of visuospatial attention. *Cerebral Cortex*, *17*(7), 1701–1711. https://doi.org/10.1093/cercor/bhl081

Irwin, D. E., Colcombe, A. M., Kramer, A. F., & Hahn, S. (2000). Attentional and oculomotor capture by onset, luminance and color singletons. *Vision Research*, *40*(10–12), 1443–1458. https://doi.org/10.1016/S0042-6989(00)00030-4

Jiang, Y., & Kumar, A. (2004). Visual short-term memory for two sequential arrays: one integrated representation or two separate representations? *Psychonomic Bulletin & Review*, *11*(3), 495–500. https://doi.org/10.3758/BF03196601

Jiang, Y., Olson, I. R., & Chun, M. M. (2000). Organization of visual short-term memory. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *26*(3), 683–702. https://doi.org/10.1037/0278-7393.26.3.683

Juan, C.-H. H., & Walsh, V. (2003). Feedback to V1: A reverse hierarchy in vision. *Experimental Brain Research*, *150*(2), 259–263. https://doi.org/10.1007/s00221-003-1478-5

Juni, M. Z., Singh, M., & Maloney, L. T. (2010). Robust visual estimation as source separation. *Journal of Vision*, *10*(14), 1–20. https://doi.org/10.1167/10.14.1

Kahrimanovic, M., Bergmann Tiest, W. M., & Kappers, A. M. L. (2009). Context effects in haptic perception of roughness. *Experimental Brain Research*, *194*(2), 287–297. https://doi.org/10.1007/s00221-008-1697-x

Kappers, A. M. L., & Bergmann Tiest, W. M. (2014). Tactile and Haptic Perceptual Organization, (November), 1–17. https://doi.org/10.1093/oxfordhb/9780199686858.013.002

Kerzel, D., & Schönhammer, J. (2013). Salient stimuli capture attention and action. *Attention, Perception, and Psychophysics*, *75*(8), 1633–1643. https://doi.org/10.3758/s13414-013-0512-3

Kietzman, M. L., & Sutton, S. (1968). The interpretation of two-pulse measures of temporal resolution in vision. *Vision Research*, *8*(3), 287–302. https://doi.org/10.1016/0042-6989(68)90016-3

Kimchi, R. (2015). The perception of hierarchical structure. *Oxford Handbook of Perceptual Organization*, (November), 129–149. https://doi.org/10.1093/oxfordhb/9780199686858.013.025

Kizuk, Sayeed, A. D., & Mathewson, K. E. (2017). Power and Phase Alpha Oscillations Reveal an Interaction between Spatial and Temporal Visual Attention. *Journal of Cognitive Neuroscience*, *29*(3), 480–494.

Koelewijn, T. K., Van Der Burg, E., Bronkhorst, A., & Theeuwes, J. (2008). Priming T2 in a visual and auditory attentional blink task. *Perception & Psychophysics*, *70*(4), 658–666.

Kool, W., Conway, A. R. A., & Turk-Browne, N. B. (2014). Sequential dynamics in visual short-term memory. *Attention, Perception, & Psychophysics*, *76*(7), 1885–1901. https://doi.org/10.3758/s13414-014-0755-7

Koshino, H. (2017). Effects of working memory contents and perceptual load on distractor processing: When a response-related distractor is held in working memory. *Acta Psychologica*, *172*, 19–25. https://doi.org/10.1016/j.actpsy.2016.11.003

Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, *14*(7), 301–307. https://doi.org/10.1016/j.tics.2010.04.006

Lamme, V. A. F. (2001). Blindsight: The role of feedforward and feedback corticocortical connections. *Acta Psychologica*, *107*(1–3), 209–228. https://doi.org/10.1016/S0001-6918(01)00020-8

Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, *23*, 571–579.

Lecerf, T., & De Ribaupierre, A. (2005). Recognition in a visuospatial memory task: The effect of presentation. *European Journal of Cognitive Psychology*, *17*(1), 47–75. https://doi.org/10.1080/09541440340000420

Lederman, S. J. (1983). Tactual roughness perception: Spatial and temporal determinants. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, *37*(4), 498–511. https://doi.org/10.1037/h0080750

Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, *77*(1), 24–42. https://doi.org/10.1152/jn.1997.77.1.24

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *NATURE*, *390*, 279–281.

Lunau, R., & Olivers, C. N. L. (2010). The attentional blink and lag 1 sparing are nonspatial. *Attention, Perception & Psychophysics*, *72*(2), 317–325.

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience*, *17*(3), 347–356. https://doi.org/10.1038/nn.3655

Macdonald, J. S. P., Cavanagh, P., & VanRullen, R. (2014). Attentional sampling of multiple wagon wheels. *Attention, Perception, and Psychophysics*, *76*(1), 64–72. https://doi.org/10.3758/s13414-013-0555-5

Macdonald, J. S. P., Mathan, S., & Yeung, N. (2011). Trial-by-trial variations in subjective attentional state are reflected in ongoing prestimulus EEG alpha oscillations. *Frontiers in Psychology*, *2*(MAY), 1–16. https://doi.org/10.3389/fpsyg.2011.00082

Mack, A., Erol, M., Clarke, J., & Bert, J. (2016). No iconic memory without attention. *Consciousness and Cognition*, *40*, 1–8. https://doi.org/10.1016/j.concog.2015.12.006

Marchant, A. P., Simons, D. J., & de Fockert, J. W. (2013). Ensemble representations: Effects of set size and item heterogeneity on average size perception. *Acta Psychologica*, *142*(2), 245–250. https://doi.org/10.1016/j.actpsy.2012.11.002

Miconi, T., & VanRullen, R. (2016). A Feedback Model of Attention Explains the Diverse Effects of Attention on Neural Firing Rates and Receptive Field Structure. *PLoS Computational Biology*, *12*(2), 1–18. https://doi.org/10.1371/journal.pcbi.1004770

Moore, C M, & Egeth, H. (1997). Perception without attention: evidence of grouping under conditions of inattention. *Journal of Experimental Psychology: Human Perception and Performance*, *23*(2), 339–352. https://doi.org/10.1037/0096-1523.23.2.339

Moore, Cathleen M. (2001). Inattentional blindness: Perception or memory and what does it matter? *Psyche*, *7*(2), 7-.

Moores, E., & Maxwell, J. P. (2008). The role of prior exposure in the capture of attention by items in working memory. *Visual Cognition*, *16*(5), 675–695. https://doi.org/10.1080/13506280701229262

Moreland, J. C., & Boynton, G. M. (2017). A neurophysiological explanation for biases in visual localization. *Attention, Perception, {&} Psychophysics*, *79*(2), 553–562. https://doi.org/10.3758/s13414-016-1251-z

Motoyoshi, I., Nishida, S., Sharan, L., & Adelson, E. H. (2007). Image statistics and the perception of surface qualities. *Nature*, *447*(7141), 206–209. https://doi.org/10.1038/nature05724

Mutluturk, A., & Boduroglu, A. (2014). Effects of spatial configurations on the resolution of spatial working memory. *Attention, Perception & Psychophysics*, 2276–2285. https://doi.org/10.3758/s13414-014-0713-4

Ngiam, W. X. Q., Khaw, K. L. C., Holcombe, A. O., & Goodbourn, P. T. (2018). Visual Working Memory for Letters Varies With Familiarity but Not Complexity. *Journal of Experimental Psychology: Learning Memory and Cognition*, *45*(10), 1761–1775. https://doi.org/10.1037/xlm0000682

Nie, Q. Y., Muller, H. J., & Conci, M. (2017). Hierarchical organization in visual working memory: From global ensemble to individual object structure. *Cognition*, *159*, 85–96. https://doi.org/10.1016/j.cognition.2016.11.009

Nieuwenhuis, S., & de Kleijn, R. (2011). Consciousness of targets during the attentional blink: A gradual or all-or-none dimension? *Attention, Perception, and Psychophysics*, *73*(2), 364–373. https://doi.org/10.3758/s13414-010-0026-1

Nieuwenstein, M. R., Chun, M. M., Van Der Lubbe, R. H. J., & Hooge, I. T. C. (2005). Delayed attentional engagement in the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(6), 1463–1475. https://doi.org/10.1037/0096-1523.31.6.1463

Nieuwenstein, M. R., Johnson, A., Kanai, R., & Martens, S. (2007). Cross-task repetition amnesia: Impaired recall of RSVP targets held in memory for a secondary task. *Acta Psychologica*, *125*(3), 319–333. https://doi.org/10.1016/j.actpsy.2006.08.006

Norman, L. J., Heywood, C. A., & Kentridge, R. W. (2015). Direct encoding of orientation variance in the visual system. *Journal of Vision*, *15*(2015), 1–14. https://doi.org/10.1167/15.4.3.doi

Olivers, C. N. L. (2008). Interactions between visual working memory and visual attention. *Frontiers in Bioscience*, *13*(3), 1182–1191. https://doi.org/10.2741/2754

Olivers, C. N. L. (2009). What drives memory-driven attentional capture? The effects of memory type, display type, and search type. *Journal of Experimental Psychology. Human Perception and Performance*, *35*(5), 1275–1291. https://doi.org/10.1037/a0013896

Olivers, C. N. L., & Eimer, M. (2011). On the difference between working memory and attentional set. *Neuropsychologia*, *49*(6), 1553–1558. https://doi.org/10.1016/j.neuropsychologia.2010.11.033

Olivers, C. N. L., & Meeter, M. (2008). A boost and bounce theory of temporal attention. *Psychological Review*, *115*(4), 836–863. https://doi.org/10.1037/a0013395

Olivers, C. N. L., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: visual working memory content affects visual attention. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(5), 1243–1265. https://doi.org/10.1037/0096-1523.32.5.1243

Olivers, C. N. L., & Nieuwenhuis, S. (2006). The Beneficial Effects of Additional Task Load , Positive Affect , and Instruction on the Attentional Blink, *32*(2), 364–379. https://doi.org/10.1037/0096-1523.32.2.364

Olivers, C. N. L., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences*, *15*(7), 327–334. https://doi.org/10.1016/j.tics.2011.05.004

Olivers, C. N. L., & Schreij, D. (2014). Visual memory performance for color depends on spatiotemporal context. *Attention, Perception & Psychophysics*, 1873–1884. https://doi.org/10.3758/s13414-014-0741-0

Papenmeier, F., Huff, M., & Schwan, S. (2012). Representation of dynamic spatial configurations in visual short-term memory. *Attention, Perception, & Psychophysics*, *74*, 397–415. https://doi.org/10.3758/s13414-011-0242-3

Pashler, H., & Shiu, L. (1999). Do images involuntarily trigger search? A test of Pillsbury's hypothesis. *Psychonomic Bulletin & Review*, *6*(3), 445–448. https://doi.org/10.3758/BF03210833

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., … Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203. https://doi.org/10.3758/s13428-018-01193-y

Peters, J. C., Goebel, R., & Roelfsema, P. R. (2009). Remembered but Unused: The Accessory Items in Working Memory that Do Not Guide Attention. *Journal of Cognitive Neuroscience*, *21*(6), 1081–1091. https://doi.org/10.1162/jocn.2009.21083

Pinto, Y., Otten, M., Seth, A., Vandenbroucke, A., Sligte, I. G., & Lamme, V. A. F. (2016). Conscious Visual Memory With Minimal Attention, *146*(2), 214–226. https://doi.org/dx.doi.org/10.1037/str0000014

Posner, M. (2016). ORIENTING OF ATTENTION: THEN AND NOW. *Q J Exp Psychol*, *69*(10), 1864–1875. https://doi.org/10.1016/j.physbeh.2017.03.040

Posner, M. I. (1980). Orienting of attention. *The Quarterly Journal of Experimental Psychology*, *32*(1), 3–25. https://doi.org/10.1080/00335558008248231

Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, *109*(2), 160–174. https://doi.org/10.1037/0096-3445.109.2.160

Pritchett, D., Gallace, A., & Spence, C. (2011). Implicit processing of tactile information: Evidence from the tactile change detection paradigm. *Consciousness and Cognition*, *20*(3), 534–546. https://doi.org/10.1016/j.concog.2011.02.006

Raab, D. H. (1963). Backward masking. *Psychological Bulletin*, *60*(2), 118–129.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, *18*(3), 849–860. https://doi.org/10.1037/0096-1523.18.3.849

Rees, G. (2007). Neural correlates of the contents of visual awareness in humans. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1481), 877–886. https://doi.org/10.1098/rstb.2007.2094

Ricker, T. J., & Cowan, N. (2015). Differences between Presentation Methods in Working Memory Procedures: A Matter of Working Memory Consolidation. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *40*(2), 417–428. https://doi.org/10.1037/a0034301.Differences

Roberts, R. D. (2013). Roughness perception across the hands. *Attention, Perception, and Psychophysics*, *75*(6), 1306–1317. https://doi.org/10.3758/s13414-013-0465-6

Roberts, R. D., & Humphreys, G. W. (2010). The role of somatotopy and body posture in the integration of texture across the fingers. *Psychological Science*, *21*(4), 476–483. https://doi.org/10.1177/0956797610363405

Robitaille, N., & Harris, I. M. (2011). When more is less: Extraction of summary statistics benefits from larger sets. *Journal of Vision*, *11*(12), 1–8. https://doi.org/10.1167/11.12.18

Rohenkohl, G., & Nobre, A. C. (2011). Alpha Oscillations Related to Anticipatory Attention Follow Temporal Expectations. *Journal of Neuroscience*, *31*(40), 14076–14084. https://doi.org/10.1523/JNEUROSCI.3387-11.2011

Rohenkohl, Gustavo, Coull, J. T., & Nobre, A. C. (2011). Behavioural dissociation between exogenous and endogenous temporal orienting of attention. *PLoS ONE*, *6*(1), 1–5. https://doi.org/10.1371/journal.pone.0014620

Rosenholtz, R. (2014). Texture Perception. *Oxford Handbook of Perceptual Organization*, (Gibson 1986), 1–24. https://doi.org/10.1016/B978-012375731-9/50081-1

Rouder, J. N., Morey, R. D., Morey, C. C., & Cowan, N. (2011). How to measure working memory capacity in the change detection paradigm. *Psychonomic Bulletin and Review*, *18*(2), 324–330. https://doi.org/10.3758/s13423-011-0055-3

Ryan, J. D., & Villate, C. (2009). Building visual representations: The binding of relative spatial relations across time. *Visual Cognition*, *17*(1–2), 254–272. https://doi.org/10.1080/13506280802336362

Saiki, J., & Holcombe, A. O. (2012). Blindness to a simultaneous change of all elements in a scene, unless there is a change in summary statistics. *Journal of Vision*, *12*(3), 2–2. https://doi.org/10.1167/12.3.2

Samaha, J., & Postle, B. R. (2015). The Speed of Alpha-Band Oscillations Predicts the Temporal Resolution of Visual Perception. *Current Biology*, *25*(22), 2985–2990. https://doi.org/10.1016/j.cub.2015.10.007

Sanocki, T., Sellers, E., Mittelstadt, J., & Sulman, N. (2010). How high is visual short-term memory capacity for object layout? *Attention, Perception & Psychophysics*, *72*(4), 1097–1109.

Sasin, E., & Nieuwenstein, M. (2016). Memory-driven attentional capture reveals the waxing and waning of working memory activation due to dual-task interference. *Psychonomic Bulletin and Review*, *23*(6), 1891–1897. https://doi.org/10.3758/s13423-016-1041-6

Schurger, A., Sarigiannidis, I., Naccache, L., Sitt, J. D., & Dehaene, S. (2015). Cortical activity is more stable when sensory stimuli are consciously perceived. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(16), E2083–E2092. https://doi.org/10.1073/pnas.1418730112

Schwartz, O., & Giraldo, L. G. S. (2017). Behavioral and neural constraints on hierarchical representations. *Journal of Vision*, *17*(3), 13. https://doi.org/10.1167/17.3.13

Sekuler, A. B., & Bennett, P. J. (2001). Generalized Common Fate: Grouping by Common Luminance Changes. *Psychological Science*, *12*(6), 437–444. https://doi.org/10.1111/1467-9280.00382

Semmelmann, K., & Weigelt, S. (2017). Online psychophysics: reaction time effects in cognitive experiments. *Behavior Research Methods*, *49*(4), 1241–1260. https://doi.org/10.3758/s13428-016-0783-4

Sengpiel, F., & Hübener, M. (1999). Visual attention: Spotlight on the primary visual cortex. *Current Biology*, *9*(9), 318–321. https://doi.org/10.1016/S0960-9822(99)80202-4

Sergent, C., & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science*, *15*(11), 720–728. https://doi.org/10.1111/j.0956-7976.2004.00748.x

Shapiro, K. L., Hanslmayr, S., Enns, J. T., & Lleras, A. (2017). Alpha, beta: The rhythm of the attentional blink. *Psychonomic Bulletin and Review*, *24*(6), 1862–1869. https://doi.org/10.3758/s13423-017-1257-0

Shin, E., Fabiani, M., & Gratton, G. (2006). Multiple levels of stimulus representation in visual working memory. *Journal of Cognitive Neuroscience*, *18*(5), 844–858. https://doi.org/10.1162/jocn.2006.18.5.844

Silvis, J. D., & Shapiro, K. L. (2014). Spatiotemporal configuration of memory arrays as a component of VWM representations. *Visual Cognition*, *22*(7), 948–962. https://doi.org/10.1080/13506285.2014.937789

Sligte, Ilga G., Scholte, H. S., & Lamme, V. A. F. (2008). Are there multiple visual short-term memory stores? *PLoS ONE*, *3*(2), 2–10. https://doi.org/10.1371/journal.pone.0001699

Sligte, Ilja G., Wokke, M. E., Tesselaar, J. P., Steven Scholte, H., & Lamme, V. A. F. (2011). Magnetic stimulation of the dorsolateral prefrontal cortex dissociates fragile visual short-term memory from visual working memory. *Neuropsychologia*, *49*(6), 1578–1588. https://doi.org/10.1016/j.neuropsychologia.2010.12.010

Slotnick, S. D. (2018). Several studies with significant C1 attention effects survive critical analysis. *Cognitive Neuroscience*, *9*(1–2), 80–90. https://doi.org/10.1080/17588928.2017.1396971

Solomon, J. A., May, K. A., & Tyler, C. W. (2016). Inefficiency of orientation averaging: Evidence for hybrid serial/parallel temporal integration. *Journal of Vision*, *16*(1), 13. https://doi.org/10.1167/16.1.13

Solomon, J. a. (2010). Visual discrimination of orientation statistics in crowded and uncrowded arrays. *Journal of Vision*, *10*(14), 19. https://doi.org/10.1167/10.14.19

Soto, D., Heinke, D., Humphreys, G. W., & Blanco, M. J. (2005). Early, involuntary top-down guidance of attention from working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(2), 248–261. https://doi.org/10.1037/0096-1523.31.2.248

Soto, D., Hodsoll, J., Rotshtein, P., & Humphreys, G. W. (2008). Automatic guidance of attention from working memory. *Trends in Cognitive Sciences*, *12*(9), 342–348. https://doi.org/10.1016/j.tics.2008.05.007

Soto, D., & Humphreys, G. W. (2007). Automatic guidance of attention from working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(3), 730–757.

Soto, D., & Humphreys, G. W. (2008). Stressing the mind: the effect of cognitive load and articulatory suppression on attentional guidance from working memory. *Perception & Psychophysics*, *70*(5), 924–934. https://doi.org/10.3758/PP.70.5.924

Soto, D., Humphreys, G. W., & Heinke, D. (2006). Working memory can guide pop-out search. *Vision Research*, *46*(6–7), 1010–1018. https://doi.org/10.1016/j.visres.2005.09.008

Soto, D., Mäntylä, T., & Silvanto, J. (2011). Working memory without consciousness. *Current Biology*, *21*(22), R912–R913. https://doi.org/10.1016/j.cub.2011.09.049

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, and Psychophysics*, *78*(7), 1839–1860. https://doi.org/10.3758/s13414-016-1108-5

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, *74*(11), 1–29.

Suchow, J. W., Fougnie, D., Brady, T. F., & Alvarez, G. a. (2014). Terms of the debate on the format and structure of visual memory. *Attention, Perception & Psychophysics*, 1–9. https://doi.org/10.3758/s13414-014-0690-7

Toh, Y. N., Sisk, C. A., & Jiang, Y. V. (2020). Effects of changing object identity on location working memory. *Attention, Perception, and Psychophysics*, *82*(1), 294–311. https://doi.org/10.3758/s13414-019-01738-z

Tong, F. (2013). Imagery and visual working memory: One and the same? *Trends in Cognitive Sciences*, *17*(10), 489–490. https://doi.org/10.1016/j.tics.2013.08.005

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

Treisman, A., & Zhang, W. (2006). Location and binding in visual working memory. *Memory and Cognition*, *34*(8), 1704–1719. https://doi.org/10.3758/BF03195932

Trick, L. M., & Pylyshyn, Z. W. (1994). Why are small and large numbers enumerated differently? A limited-capacity preattentive stage in vision. *Psychological Review*, *101*(1), 80–102. https://doi.org/10.1037/0033-295X.101.1.80

Tsubomi, H., Fukuda, K., Watanabe, K., & Vogel, E. K. (2013). Neural Limits to Representing Objects Still within View. *Journal of Neuroscience*, *33*(19), 8257–8263. https://doi.org/10.1523/JNEUROSCI.5348-12.2013

Udale, R., Farrell, S., & Kent, C. (2017). No evidence for binding of items to task-irrelevant backgrounds in visual working memory. *Memory & Cognition*, 1144–1159. https://doi.org/10.3758/s13421-017-0727-y

Uttal, W. R., Spillmann, L., Stürzel, F., & Sekuler, A. B. (2000). Motion and shape in common fate. *Vision Research*, *40*(3), 301–310. https://doi.org/10.1016/S0042-6989(99)00177-7

Van Der Horst, B. J., Duijndam, M. J. A., Ketels, M. F. M., Wilbers, M. T. J. M., Zwijsen, S. A., & Kappers, A. M. L. (2008). Intramanual and intermanual transfer of the curvature aftereffect. *Experimental Brain Research*, *187*(3), 491–496. https://doi.org/10.1007/s00221-008-1390-0

van Moorselaar, D., Theeuwes, J., & Olivers, C. N. L. (2014). In competition for the attentional template: Can multiple items within visual working memory guide attention? *Journal of Experimental Psychology: Human Perception and Performance*, *40*(4), 1450–1464. https://doi.org/10.1037/a0036229

Vandenbroucke, A. R. E., Sligte, I. G., & Lamme, V. A. F. (2011). Manipulations of attention dissociate fragile visual short-term memory from visual working memory. *Neuropsychologia*, *49*(6), 1559–1568. https://doi.org/10.1016/J.NEUROPSYCHOLOGIA.2010.12.044

VanRullen, R. (2016). Perceptual Cycles. *Trends in Cognitive Sciences*, *20*(10), 723–735. https://doi.org/10.1016/j.tics.2016.07.006

Vanrullen, R., & MacDonald, J. S. P. (2012). Perceptual echoes at 10 Hz in the human brain. *Current Biology*, *22*(11), 995–999. https://doi.org/10.1016/j.cub.2012.03.050

Vidal, J. R., Gauchou, H. L., Tallon-Baudry, C., & O'Regan, J. K. (2005). Relational information in visual short-term memory: the structural gist. *Journal of Vision*, *5*(3), 244–256. https://doi.org/10:1167/5.3.8

Vogel, E. K., Woodman, G. F., & Luck, S. J. (2006). The time course of consolidation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(6), 1436–1451. https://doi.org/10.1037/0096-1523.32.6.1436

Ward, E. J., Bear, A., & Scholl, B. J. (2016). Can you perceive ensembles without perceiving individuals?: The role of statistical perception in determining whether awareness overflows access. *Cognition*, *152*, 78–86. https://doi.org/10.1016/j.cognition.2016.01.010

Wolfe, J. M. (1994). Guided Search 2 . 0 A revised model of visual search. *Psychonomic Bulletin & Review*, *1*(2), 202–238.

Woodman, G F, Vogel, E. K., & Luck, S. J. (2001). Visual search remains efficient when visual working memory is full. *Psychological Science*, *12*(3), 219–224. https://doi.org/10.1111/1467-9280.00339

Woodman, Geoffrey F, & Luck, S. J. (2007). Do the contents of visual working memory automatically influence attentional selection during visual search? *Journal of Experimental Psychology. Human Perception and Performance*, *33*(2), 363–377. https://doi.org/10.1037/0096-1523.33.2.363

Woodman, Geoffrey F, Vogel, E. K., & Luck, S. J. (2012). Flexibility in Visual Working Memory: Accurate Change Detection in the Face of Irrelevant Variations in Position. *Visual Cognition*, *20*(1), 1–28. https://doi.org/10.1080/13506285.2011.630694

Wyart, V., & Tallon-Baudry, C. (2008). Neural Dissociation between Visual Awareness and Spatial Attention. *Journal of Neuroscience*, *28*(10), 2667–2679. https://doi.org/10.1523/JNEUROSCI.4748-07.2008

Wyble, B., Potter, M. C., Bowman, H., & Nieuwenstein, M. (2011). Attentional episodes in visual perception. *Journal of Experimental Psychology: General*, *140*(3), 488–505. https://doi.org/10.1037/a0023612

Yantis, S. (2005). How visual salience wins the battle for awareness. *Nature Neuroscience*, *8*(8), 975–977. https://doi.org/10.1038/nn0805-975

Yantis, S., & Jonides, J. (1988). Uniqueness of abrupt visual onset in capturing attention. *Perception & Psychophysics*, *43*(4), 346–354.

Yoshida, T., Yamaguchi, A., Tsutsui, H., & Wake, T. (2015). Tactile search for change has less memory than visual search for change. *Attention, Perception, and Psychophysics*, *77*(4), 1200–1211. https://doi.org/10.3758/s13414-014-0829-6

Yoshioka, T., Craig, J. C., Beck, G. C., & Hsiao, S. S. (2011). Perceptual constancy of texture roughness in the tactile system. *Journal of Neuroscience*, *31*(48), 17603–17611. https://doi.org/10.1523/JNEUROSCI.3907-11.2011

Zauner, A., Fellinger, R., Gross, J., Hanslmayr, S., Shapiro, K., Gruber, W., … Klimesch, W.

(2012). Alpha entrainment is responsible for the attentional blink phenomenon. *NeuroImage*, *63*(2), 674–686. https://doi.org/10.1016/j.neuroimage.2012.06.075

Zhang, B., Zhang, J. X., Huang, S., Kong, L., & Wang, S. (2011). Effects of load on the guidance of visual attention from working memory. *Vision Research*, *51*(23–24), 2356–2361. https://doi.org/10.1016/j.visres.2011.09.008

Zhang, B., Zhang, J. X., Kong, L., Huang, S., Yue, Z., & Wang, S. (2010). Guidance of visual attention from working memory contents depends on stimulus attributes. *Neuroscience Letters*, *486*(3), 202–206. https://doi.org/10.1016/j.neulet.2010.09.052

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), 233–235. https://doi.org/10.1038/nature06860

Zhao, L., Gao, Q., Ye, Y., Zhou, J., Shui, R., & Shen, M. (2014). The role of spatial configuration in multiple identity tracking. *PLoS ONE*, *9*(4). https://doi.org/10.1371/journal.pone.0093835

Zimmermann, E., Schnier, F., & Lappe, M. (2010). The contribution of scene context on change detection performance. *Vision Research*, *50*(20), 2062–2068. https://doi.org/10.1016/j.visres.2010.07.019