Biases in Information Processing Across Perception and Working Memory

Theresa Wildegger

St John's College

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Trinity Term 2016

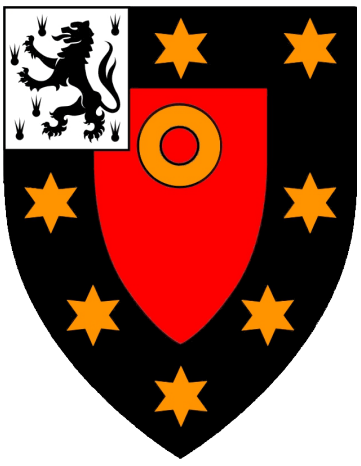

Abstract

Biases in Information Processing Across Perception and Working Memory

Theresa Wildegger

St John's College, Trinity Term 2016

A dissertation submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford.

Research over the last decades has demonstrated an enormous flexibility of attentional mechanisms in the human brain and has revealed many forms of biases across perception and working memory. In this thesis, I extend our knowledge of the flexible role of attention, and ways in which information processing can be biased to guide perception and working memory (WM). In the General Introduction, I review perceptual, WM and attention literature to describe our current understanding of the many biases across different stages of information processing. In Chapter 2, I present a study that uses functional magnetic resonance imaging (fMRI) and magnetoencephalography (MEG) to contrast and compare neural mechanisms supporting spatial and feature-based attentional selection. Neural signatures of preparatory attention revealed that comparable neural mechanisms support both kinds of attention. However, brain activity is modulated relatively more globally for feature-based compared to spatial attention, consistent with representational differences between feature and spatial information. In Chapter 3, I describe four Experiments exploring the impact of incoming sensory information, presented supraliminally or subliminally, on performance in a WM task and a visuomotor priming task. While both subliminal and supraliminal information shaped performance in the visuomotor priming task, WM representations appeared to be protected against subliminal influences and performance was only affected by information presented supraliminally. In Chapter 4, I present three Experiments that examine how bottom-up stimulus differences during encoding modulate top-down attentional effects in WM. The results reveal that benefits of retrospective attentional selection in WM persist regardless of the quality of sensory information, but the underlying mechanisms vary along with changes in stimulus properties during encoding. In Chapter 5, the General Discussion, I place the results of this work in their wider context, and discuss limitations and future directions.

Approximate Word Count: 50,000 words

Funded by the Economic and Social Research Council and St John's College, Oxford.

Dedicated to the memory
of my late supervisor, Professor Glyn Humphreys.

Acknowledgements

Perhaps this is the section I struggled with the most. There are countless people that have defined my experiences throughout my time at Oxford and the course of my DPhil. There is no doubt that a single paragraph for each ‘category’ will not do their support justice, or my appreciation thereof. My hope is that every single one knows how thankful I am for their help and friendship.

Both my supervisors, Kia Nobre and Glyn Humphreys, have given me invaluable guidance and support. I consider myself lucky to have worked with such inspirational, intelligent and compassionate individuals. They have shown me how to keep going after fallbacks, the beauty of facing challenges, and they always (somehow) found time for my concerns and questions. Because of them, I have grown as an academic as well as a person.

I am also very grateful to Nick Myers, Freek van Ede, Celine Gillebert, Nahid Zokaei, Zita Patai, Robert Mok, Ian Gould, George Wallis, Gustavo Rohenkohl, Nele Demeyere and Mihaela Duta. Without you, this thesis would look very different, and so would the contents of my brain.

Thanks also to Matt, Niina, Nick, Gabi (the human), Petya, Josh, and Melissa for emotional support throughout this thesis – be it in the form of pints in the pub, runs in the park, shared complaints in the departmental hallway, or just an open ear and a smile & nod.

A special thanks to my college advisor, Kate Nation, for support throughout my DPhil and MSc, and for helping me navigate through academic life and many challenges along the way.

Finally, a big THANK YOU to my mum, my dad, my brothers, my grandparents, Petra and Gabe. You all know why.

Table of Contents

1 General Introduction

1.1 Overview

Our senses are constantly bombarded with vast amounts of data while our capacity to process this onslaught is strictly limited. Attention mitigates this constraint by biasing the allocation of processing resources in favour of task-relevant information. Classic questions in attention concern this 'bottleneck' of information processing: How is incoming information selected for further processing to optimise goal-directed behaviour (e.g. Broadbent, 1958; Deutsch & Deutsch, 1963)? What is the unit of attentional selection? At what stage of processing is information selected, and what is the fate of information that has not been selected? Early research often seemed to demonstrate that attentional mechanisms primarily influence specific stages or representations in information processing. For example, early-selection theories (Broadbent, 1958) argued all attentional selection occurs at perceptual stages in the processing stream. In contrast, late-selection theories reasoned all selection happens after identification has taken place (Deutsch & Deutsch, 1963). Similarly, research concerning attentional reference frames, i.e. its representational basis, favoured either object-based (Baylis & Driver, 1993; Duncan, 1984; Rock & Gutman, 1981) or space-based (Broadbent, 1982; Posner, 1980) attentional selection.

However, decades of research have revealed an enormous flexibility of attentional mechanisms. Attention can operate on individual features, objects, locations in space, and time; on perceptual templates or on working memory (WM) templates (Baylis & Driver, 1992; Duncan, 1984; Egly, Driver, & Rafal, 1994; Griffin & Nobre, 2003; Harms & Bundesen, 1983; Martinez et

al., 2006; Maunsell & Treue, 2006; Posner, 1980; Posner, Cohen, & Rafal, 1982; Rerko, Souza, & Oberauer, 2014; Vecera & Farah, 1994; Williams, Hong, Kang, Carlisle, & Woodman, 2013). This short list provides just a broad overview and any of these items could be further differentiated. For example, attention mechanisms are involved in all stages of WM processing: the encoding of information into WM, maintenance of information in WM, and efficient retrieval of WM items (see Gazzaley & Nobre, 2012 for a review). The term attention has come to be used to describe a range of phenomena related to biased selection and integration of information that is relevant to current behavioural goals.

In addition to attentional mechanisms, research has uncovered a wealth of other biases in information processing. For example, top-down memory states are a source of biasing signals, and WM contents can guide attention and bias perception and conscious experience. When an item is held in WM, physical stimuli that match the WM item are more likely to capture attention and enter awareness (Gayet, Paffen, & Van der Stigchel, 2013a; Soto, Hodsoll, Rotshtein, & Humphreys, 2008). Furthermore, internal representations can be biased by perceptual stimuli presented before, during, or after encoding has taken place. For example, past (sensory) experiences modulate responses to subsequent stimuli: repeated presentation of the same stimulus facilitates perception of, and motor responses to that stimulus, a phenomenon known as priming (see Tulving & Schachter, 2013 for a review). Beyond facilitation, previously encountered stimuli can also change perception and WM representations of subsequently-presented stimuli (Dubé, Zhou, Kahana, & Sekuler, 2014; Fischer & Whitney, 2014; Huang & Sekuler, 2010). Similarly, the presence

of task-irrelevant stimuli during stimulus encoding changes perceptual responses to the target stimulus (e.g. Gilbert & Wiesel, 1990). Once task-relevant information has been encoded and is no longer present in the environment, an existing internal representation can be influenced by incoming, sensory information, even if this new information is task-irrelevant (Magnussen, 2000). Finally, spontaneous fluctuations in baseline levels of neural activity also exert an influence on how incoming information is processed and encoded into WM (e.g. Gutnisky, Beaman, Lew, & Dragoi, 2016; Mathewson, Gratton, Fabiani, Beck, & Ro, 2009; Myers, Stokes, Walther, & Nobre, 2014). In addition to top-down goal states, top-down memory states, as well as internal (neural) and external (physical) contexts shape how information is processed.

The experimental work presented in this thesis addresses the extent and characteristics of these influences. First, I review studies showing how attentional mechanisms can flexibly operate on the basis of feature or spatial information to modulate perceptual processes (Behrmann & Tipper, 1999; Egly et al., 1994; Posner, 1980). Attention can be directed not only to regions in space, but also to overlapping features that cannot be separated by location alone. While the effects of both feature-based attention and spatial attention on perceptual processing are relatively well researched in isolation, almost no research has directly compared feature-based attention and spatial attention (but see Giesbrecht, Woldorff, Song, & Mangun, 2003; McAdams & Maunsell, 2000), a significant gap that will be addressed in this thesis. I discuss the differences in the neural representation of feature and spatial information and then review behavioural, neuroimaging and electrophysiological studies showing striking similarities between the effects

of feature-based attention and spatial attention. While have a relatively clear idea about what kind of mechanisms could support spatial attention, it is much less clear how attentional mechanisms support feature-based selection given the different representation of feature information. Distinct locations are represented in different brain areas, but distinct features are represented in overlapping brain areas and are, therefore, more difficult to bias. In Chapter 2 of my thesis I compare and contrast how feature-based and spatial information can be used to guide attention and optimise perception.

Following on from this, I discuss interactions between perceptual and WM processes. I review evidence that contents of WM guide selective attention and bias perception (Desimone & Duncan, 1995; Downing, 2000; Soto et al., 2008). I then move on to discuss whether such influences also occur in the other direction, i.e. whether and how perceptual input could bias existing memory contents. While effects of WM contents on perceptual processing are well established, this question is much less well researched. I review research on the effects of perceptual distracters on WM performance and highlight important details that remain unresolved: does distracter presentation simply lead to information loss, a change in WM representations, or both? How are any such influences modulated by signal strength[1] of the incoming (distracter) information? Are effects dependent on target-distracter similarity? What is the role of task-relevance of incoming perceptual information? Chapter 3 establishes how WM contents are

[1] I use the term 'signal strength' loosely and in relation to any (bottom-up) stimulus characteristics that might influence how much information can be retained from its presentation. This could be due to changes in exposure duration, contrast level, amount of signal (e.g. coherently moving dots in random dot motion stimuli), among other things.

shaped by the presentation of perceptual information, and describes how signal strength and task-relevance modulate the observed effects.

Finally, I review research on biasing mechanisms occurring within WM. I discuss research showing how WM contents are shaped by attentional selection. I place particular emphasis on the 'retrocueing' paradigm where observers use a cue presented during the WM delay period, but well after the time window of iconic memory, to orient attention retrospectively within WM (e.g. Griffin & Nobre, 2003; Landman, Spekreijse, & Lamme, 2003). Observers benefit from such a cue and performance is better following valid cues compared to neutral or invalid cues. I then review another large body of research that has revealed how effects of perceptual effects of attentional selection vary with the relative signal strength of the selected stimulus, in these studies varied by changing contrast levels or levels of external noise (e.g. Lu & Dosher, 1998; Martínez-Trujillo & Treue, 2002; Reynolds, Pasternak, & Desimone, 2000). I discuss how retrocueing effects might vary with the relative signal strength of a stimulus during WM encoding, a topic that has not been addressed before. Then, I discuss how interactions between items in WM might be shaped by attentional selection and by the representational quality of items in WM (as determined by stimulus strength during encoding). In Chapter 4 of my thesis, I examine how top-down attentional mechanisms affect WM representations of varying quality, and how they affect interactions between WM representations.

1.2 Feature and spatial templates guide perception

1.2.1 Differences in neural representation of feature and spatial information

Considering what is known about the neural representation of feature information, it is difficult to establish a conceptual framework for top-down feature compared to top-down spatial attention mechanisms. In the case of spatial attention, we know that space is represented retinotopically in much of the cortex, so spatial attention can target the responses of localised neurons whose receptive fields (RFs) correspond to the task-relevant location. In contrast, it is not clear how such mechanisms could also support feature-based attention. Unlike visual space, which consists of only two dimensions, there are many different non-spatial features at different levels of complexity (e.g. colours and lines versus letters and shapes). The exact details of the neural representation of feature information are unclear (Maunsell & Treue, 2006). While we lack a detailed understanding of the neural representation of features, it has been established that neurons coding for features are spatially dispersed throughout the cortex. Feature-based attention mechanisms, then, have to target spatially dispersed neurons in a feature-selective manner, while spatial attention targets neurons corresponding to the attended location, but in a non-feature-selective fashion. In other words, spatial attention could operate by modulating neighbouring neurons through a mechanism that relies on retinotopic organisation while feature-based attention has to operate more globally and requires a mechanism that does not rely on topographic organisation.

1.2.2 Similar neural mechanisms at global versus local scales

Findings from a number of studies converge to suggest that both spatial attention and feature-based attention rely on similar mechanisms but at different spatial scales (e.g. Cohen & Maunsell, 2011; Saenz, Buracas, & Boynton, 2002). For one, spatial attention and feature-based attention have comparable effects on behaviour, improving accuracy, RTs, and detection sensitivity (e.g. Carrasco, 2014; Posner, 1980; Posner, Snyder, & Davidson, 1980). Second, spatial attention and feature-based attention have comparable effects on local neural response profiles: both boost the gain of population responses and neurons that are selective to the attended location or feature respond more strongly (Martínez-Trujillo & Treue, 2002; McAdams & Maunsell, 1999; Reynolds et al., 2000; Treue & Martínez Trujillo, 1999; Williford & Maunsell, 2006). Third, similar top-down control mechanisms support spatial attention and feature-based attention (Baldauf & Desimone, 2014; Giesbrecht et al., 2003). For example, Baldauf and Desimone (2014) showed that neural mechanisms supporting non-spatial (in this case object-based) attention are similar to those supporting spatial attention. Participants were presented with temporally and spatially overlapping streams of faces and houses and asked to attend to either the face stream or the house stream while brain activity was recorded using MEG and fMRI. Attention induced oscillatory synchronisation in the gamma band (> 30 Hz) between the inferior frontal junction (IFJ), a region in the prefrontal cortex, and the fusiform gyrus or parahippocampal gyrus, when faces or houses were attended, respectively. This pattern is comparable to spatial attention, which has been shown to induce oscillatory synchrony between frontal eye fields (FEF), another region in the prefrontal cortex, and

the visual cortex, with a time shift consistent with the proposal that the directionality of the effect is from FEF to visual areas (Gregoriou, Gotts, Zhou, & Desimone, 2009a). In both spatial and non-spatial attention, regions in prefrontal cortex send a top-down signal to cortical regions that process the attended feature or location to bias neural processing.

While most studies have examined either spatial attention or feature-based attention separately, there are some notable exceptions that have directly compared spatial and feature-based attention in the same experiment (Cohen & Maunsell, 2010; Giesbrecht et al., 2003). These studies highlight similarities between the two kinds of attention, but have additionally provided strong evidence for the proposal that feature-based attention mechanisms operate globally, while spatial attention mechanisms operate at a more local level. First, Giesbrecht et al. (2003) compared feature-based to spatial attention to test the generalizability of the frontoparietal control network, a distributed network of interconnected brain regions involved in executive control functions (Corbetta, Miezin, Shulman, & Petersen, 1993; Mesulam, 1990; Nobre et al., 1997). Participants performed a cued (feature or spatial) attention task while their brain activity was recorded using fMRI. Areas of activation were largely identical between the two attention conditions, supporting the notion that the frontoparietal control network is not specific to spatial attention but is instead a general executive network that is recruited for both feature-based and spatial attention.

A more recent study compared the effects of feature-based and spatial attention on neuronal activity in monkey area V4 (Cohen & Maunsell, 2011). The effects on local populations of neurons were comparable: an

increase in firing rates to preferred stimuli and a decrease in firing rates to non-preferred stimuli. However, spatial attention targeted only local populations within one hemisphere while effects of feature-based attention were coordinated across both hemispheres. Together, these studies suggest that feature-based and spatial attention exert comparable effects on neural activity but spatial attention operates locally while feature-based attention operates more globally across the cortex. This conclusion is also consistent with psychophysical and neuroimaging studies that show global biasing effects of feature-based attention across the entire visual field (Folk, Leber, & Egeth, 2002; Jehee, Brady, & Tong, 2011; Saenz et al., 2002; Serences & Boynton, 2007).

To summarise, at the level of local neural activity, both feature-based and spatial attention increase gain in neural populations but spatial attention does so in a feature non-selective manner while feature-based attention modulates response gain in a feature-selective way. This modulation is driven by top-down signals originating from a frontoparietal network. Importantly, a spatially dispersed representation of features requires that such top-down signals can target only neurons tuned to the attended feature among neurons with different selectivity profiles. How such feature-selective gain modulation could be implemented globally, across the whole cortex, is an open question.

1.2.3 Neural oscillations and attention

One mechanism that has been proposed to underlie gain modulation is synchronized oscillatory activity (Fries, 2009). Indeed, studies have demonstrated a key role of oscillatory activity in selective attention both in

humans and monkeys. One of the earliest studies to propose a link between synchronized firing patterns and selective attention was reported by (Steinmetz et al., 2000). Three monkeys were trained to switch attention between a visual and a tactile discrimination task while perceptual input remained constant. The firing between pairs of neurons in the secondary somatosensory cortex was recorded throughout the task. The degree of synchrony changed with attentional shifts: there was a relative increase in the synchronous firing between neuron pairs when animals focused on the tactile task relative to when they focused on the visual task.

Another early study also linked oscillatory synchronisation to selective attention and assessed the contributions of high-frequency oscillations in the gamma band (> 30 Hz) and low-frequency oscillations (< 17 Hz) separately (Fries, Reynolds, Rorie, & Desimone, 2001a). The authors recorded activity from macaque area V4 while they presented monkeys with two stimuli. Animals were cued to attend to one of the two stimuli, and neuronal activity was compared between trials where the animal attended to the stimulus in the RF to trials were attention was directed to a stimulus outside the RF. High-frequency gamma synchronization was increased with attention, while low-frequency synchronization were decreased, both during the expectation period and in response to stimulus presentation. These studies demonstrate a link between attentional processing and oscillatory synchronization in low and high frequency bands.

There is now a substantial body of evidence supporting these early findings (Jensen, Kaiser, & Lachaux, 2007; Klimesch, 2012) but there has been no direct comparison between the effects of feature-based and spatial attention on high- and low-frequency oscillations. It is not clear how the

same oscillatory mechanisms could support the attentional selection on the basis of feature and spatial information, considering the extent to which they differ in their neural representation. In the case of low-frequency oscillations, research has shown that both feature-based and spatial attention modulate oscillations in the alpha band (8 – 14 Hz) during preparatory periods of attention tasks (Bollimunta, Mo, Schroeder, & Ding, 2011; Buffalo, Fries, Landman, Buschman, & Desimone, 2011; Fries, Reynolds, Rorie, & Desimone, 2001a; Fries, Womelsdorf, Oostenveld, & Desimone, 2008; Müller & Keil, 2004; Sauseng et al., 2005; Siegel, Donner, Oostenveld, Fries, & Engel, 2008; Snyder & Foxe, 2010; Thut, Nietzel, Brandt, & Pascual-Leone, 2006; Worden, Foxe, Wang, & Simpson, 2000; Wyart & Tallon-Baudry, 2008a). Alpha power reductions observed following shifts of attention are generally interpreted to reflect a release of inhibition of neurons coding for attended features or regions in space while neurons coding for irrelevant features and regions of space remain inhibited (Bauer, Stenner, Friston, & Dolan, 2014a; Klimesch, Sauseng, & Hanslmayr, 2007; but see Bauer, Stenner, Friston, & Dolan, 2014b). The finding that the topography of alpha power modulations closely corresponds to the task-relevant location or sensory modality supports this interpretation. For example, Kelly, Lalor, Reilly, & Foxe (2006) used a spatial cueing task where either the left or right hemifield was cued. Alpha power modulations were observed contralateral to the cued region of space. Snyder and Foxe (2010) used a feature-dimension-cueing task and in different trials either colour or motion were task-relevant. Alpha power modulations were observed bilaterally either along the dorsal stream in motion area MT, or along the ventral stream in colour area V4, in line with whichever feature

dimension was currently relevant. These findings suggest that alpha power modulations can be employed flexibly on the basis of both feature and spatial information, to disinhibit those neural populations that are coding task-relevant information.

Low-frequency alpha power modulations are linked to both spatial and feature-based attention but the target sites are very different between the two kinds of attention. In the case of spatial attention, neurons are biased locally in a topographic manner while for feature-based attention alpha modulations occur globally, distributed across both hemispheres. Since there are no direct comparisons between feature-based and spatial attention it is unclear whether alpha power modulations differ in their time course and implementation, or whether there are interactive effects between feature-based and spatial attention.

In the case of high-frequency oscillations, research has shown that gamma band (>30 Hz) power is increased both by feature-based and spatial attention, primarily in response to stimulus presentation (Bauer, Oostenveld, Peeters, & Fries, 2006; Bichot, Rossi, & Desimone, 2005; Buffalo et al., 2011; Fries et al., 2008; Fries, Reynolds, Rorie, & Desimone, 2001b; Gruber, Müller, Keil, & Elbert, 1999; Müller & Keil, 2004; Siegel et al., 2008; Vidal, Chaumon, O'Regan, & Tallon-Baudry, 2006; Womelsdorf, Fries, Mitra, & Desimone, 2006). The interpretation of gamma power modulations is less straightforward than that of alpha power modulations. Synchronous firing of groups of neurons in the gamma range increases their impact on postsynaptic targets because spikes are temporally focused as they arrive within a few milliseconds of each other at target sites (Fries, 2009). An increase in the oscillatory synchrony of those neurons supporting the target

representation increases their synaptic efficiency, which leads to an amplification of the target-related signal. This gives the attended item a competitive advantage throughout later stages of information processing. Recent proposals invoke a fundamental role for gamma in numerous higher cognitive functions, of which selective attention is one (Fries, 2009; Tallon-Baudry, 2009; Vidal et al., 2006). According to this proposal, information processing in general benefits from gamma synchronization. If this is the case, there is little reason to expect differential gamma power modulations for feature-based compared to spatial attention, except for a similar topographic differentiation that I described for alpha synchronization. Indeed, Baldauf and Desimone (2014) showed that non-spatial attention induced gamma power synchronization between pre-frontal regions and cortical regions specialized for the processing of the attended object, while Gregoriou et al. (2009) demonstrated that spatial attention induced gamma power synchronization between pre-frontal regions and cortical regions specialized for the processing of the attended location (as already discussed earlier).

However, other proposals have linked gamma band synchronization to a variety of other, more specialised functions, and its exact role in different aspects of attentional processing is unclear. For example, in addition to attentional selection, gamma power modulations have been linked to feature-binding and coherent object perception (Eckhorn et al., 1988; Singer & Gray, 1995; Tallon-Baudry, Bertrand, Delpuech, & Permier, 1997; Tallon-Baudry, Bertrand, Delpuech, & Pernier, 1996), visual awareness (Wyart & Tallon-Baudry, 2008b), and memory (Honkanen, Rouhinen, Wang, Palva, & Palva, 2014; Pesaran, Pezaris, Sahani, Mitra, &

Andersen, 2002; Pulvermüller, Keil, & Elbert, 1999; Tallon-Baudry, Bertrand, Peronnet, & Pernier, 1998). Of particular interest for the current question is the proposal that gamma oscillations play a unique role in maintaining feature-based information and object representations, which has emerged from several lines of research (Herrmann, Munk, & Engel, 2004; Honkanen et al., 2014; Tallon-Baudry, Kreiter, & Bertrand, 1999). For example, an increase in gamma power is observed when observers view coherent and gestalt-like stimuli compared to when the same image is viewed but no coherent object experienced[2] (Tallon, Bertrand, Bouchet, & Pernier, 1995; Tallon-Baudry et al., 1996). Gamma power also increases when observers view familiar objects compared to unfamiliar objects (Herrmann et al., 2004). On the basis of these results Honkanen et al. (2014) tested the proposal that gamma band activity is involved in the maintenance of feature-specific sensory information. Observers completed a delayed match-to-sample task for the colour, location or shape of two or four objects. After a delay period observers responded to a test stimulus with a forced-choice discrimination along the relevant feature dimension. Gamma oscillation amplitudes in the memory delay differed between memorized features and were modulated by memory load for shape and colour memory, but not for location memory. These shape- and colour- specific load effects predicted behavioural performance and were most pronounced in visual cortical regions, particularly in those areas that support the processing of those feature. The

[2] It is worth mentioning that some of these results have been criticized as scalp-recorded EEG can suffer from eye movement artifacts, and results may be driven by differences in eye movements between conditions (Yuval-Greenberg, Tomer, Keren, Nelken, & Deouell, 2008).

authors concluded that gamma oscillations reflect the maintenance of feature-specific information.

These findings suggest that gamma power modulations play a unique role in feature maintenance. By extension, this predicts a role for gamma power modulations in feature-based compared to spatial attention since feature-based attention likely requires the maintenance of a feature template. Feature-based attention might show a greater reliance on gamma power modulations, reflected in relatively more pronounced changes in gamma power compared to spatial attention. Alternatively, there might be a functional distinction within the gamma band for feature-based versus spatial attention. 'Gamma' refers to frequencies above >30 Hz and studies have reported attention-related gamma effects in many different frequency bands within the gamma range, at relatively low frequencies around 35 - 51 Hz (Gruber et al., 1999) up to higher ranges such as 60 – 95 Hz (Bauer et al., 2006). To some extent these differences are probably related to different recording techniques used, however, evidence that distinct frequency ranges within the gamma band might support distinct functions has been reported (Wyart & Tallon-Baudry, 2008b). In this study, a functional dissociation between different ranges of the gamma band was demonstrated: 76 Hz to 90 Hz was modulated by spatial attention while 54 Hz to 64 Hz was modulated by visual awareness. Perhaps different frequency bands within the gamma band support feature-based selection versus spatial selection.

In Chapter 1, I contrast feature-based and spatial attention directly in the same experiment using MEG and fMRI, and explore the relative role of gamma and alpha oscillations in these two types of selective attention.

Given the differences in the neural representation of features and space, the question remains if and how the same oscillatory mechanisms could support the attentional modulation of neural activity on the basis of feature compared to spatial information.

1.3 Interactions between WM contents and perception

According to the Biased Competition Model (Desimone & Duncan, 1995) any short-term description of the currently task-relevant information can be a source of biasing signals in information processing. That is, both attention and WM templates can exert biases on perceptual processing. Indeed, the overlap between WM and attentional mechanisms is so extensive it has been proposed that both are supported by the same top-down mechanisms (Gazzaley and Nobre, 2012; Awh and Jonides, 2001; Chun and Johnson, 2011). A topic that is much less well understood is whether bottom-up sensory input also shapes existing memory contents. Could information that is already in WM be altered by incoming perceptual information and what would be the nature of such influences?

1.3.1 Distracter interference in WM tasks

Within the WM literature, effects of distracters on task performance have been well researched. The presentation of a distracter, both before the WM item has been encoded and during the maintenance period, impairs WM performance (Bennett & Cortese, 1996; Jha, Fabian, & Aguirre, 2004; Lalonde & Chaudhuri, 2002; Magnussen, Greenlee, Asplund, & Dyrnes, 1991; Nemes, Parry, Whitaker, & McKeefry, 2012; Pasternak & Zaksas, 2003; Postle, 2005). Even when distracters are task-irrelevant, visual

cortical activity is not resistant to interference (Yoon, Curtis, & D'Esposito, 2006).

Interestingly, distracters tend to have more impact on WM performance when they are similar (but not identical) in appearance to the target item (Dolcos, Miller, Kragel, Jha, & McCarthy, 2007; Sreenivasan & Jha, 2007). However, interference effects disappear when distracter stimuli are identical to WM contents (Bennett & Cortese, 1996; Magnussen et al., 1991). This has been shown in psychophysical studies on memory for basic visual features (e.g. spatial frequency, orientation, motion, and contrast) which revealed that a distracter stimulus presented between the target and probe item, but outside the range of iconic memory, only interferes with performance when its value is different to the target feature. When the distracter feature matches the target value no interference is observed (see (Magnussen, 2000 for a review).

Magnussen (2000) has named this phenomenon memory masking, and offers an explanation where memory for a given feature dimension is supported by memory stores in the visual cortex that each code a restricted range of feature values along the feature dimension. These stores are connected with lateral inhibitory connections and distracter effects emerge when different memory stores maintain conflicting information. Such a conflict leads to information loss as memory stores inhibit one another. According to this feature-specific memory stores account, interference effects are due to information loss arising from competition between representations.

While a proposal based on information loss can account for the findings, various other studies point to an alternative account where

information is not lost, but changed. Incoming perceptual information might alter WM representations, as opposed to (or in addition to) reducing information. If distracter presentation changes the target representation away from the target value, a drop in performance would also be observed even if the target representation remains otherwise unchanged. Figure 1-1 provides a schematic of some possible scenarios that could equally well account for observed interference effects.

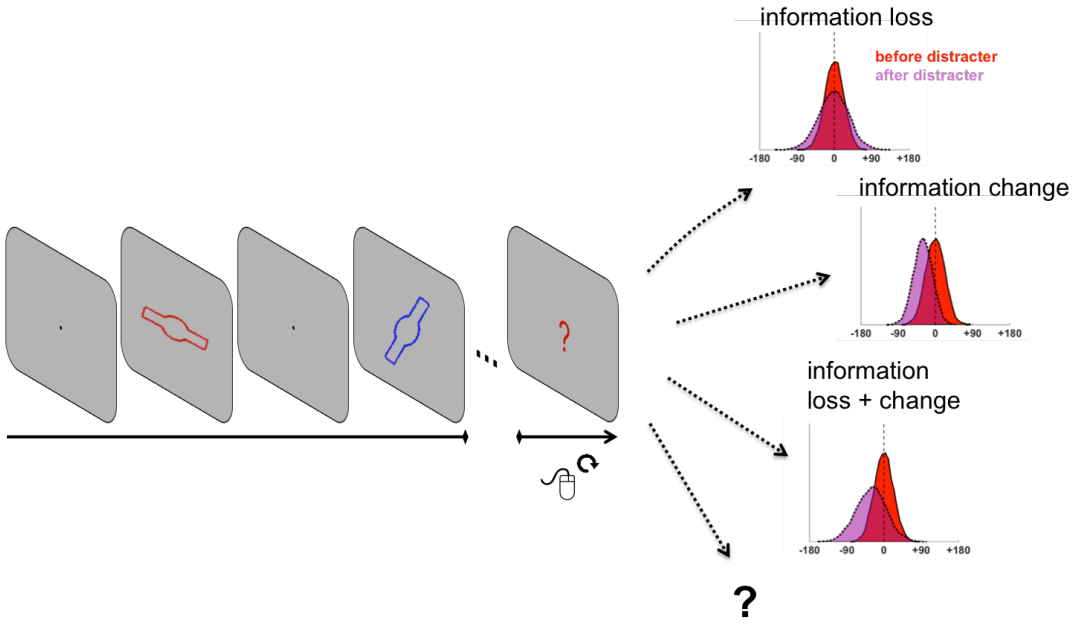

Figure 1-1. This schematic demonstrates possible effects of distracter presentation on WM representations. A distracter might lead to information loss, a shift in the target representation toward (or away) from the distracter value, or it might lead to both information loss and a shift toward (away from) the distracter value. For clarity, this schematic does not consider influences of selective attention, target-distracter similarity, or signal strength of distracter and target information, all of which are discussed in the main text.

For example, distracter effects might also be linked to serial dependence, a phenomenon in visual perception recently demonstrated by Fischer and Whitney (2014). They showed that memory traces for stimuli presented several seconds earlier bias the perception of current stimuli

toward the identity of the previously seen stimulus. This effect does not require long adaptation to previous stimuli but instead emerges after brief exposure to stimuli presented for only 500 ms. It was proposed that such an effect might emerge from neural gain changes or tuning shifts following stimulus exposure, which would affect the encoding of subsequent stimuli.

Importantly, Huang and Sekuler (2010) demonstrated that similar effects emerge in WM tasks. Observers were presented with two sequential Gabors of varying spatial frequency and had to report the spatial frequency of one of the two Gabors on a continuous sliding scale. Recall of the spatial frequency target was biased towards the spatial frequency of the distracter item. This finding could also be described as serial dependence in WM: past exposure to distracters biases memory toward the distracter value. Interestingly, Huang and Sekuler (2010) also reported a biasing effect when the distracter was presented after the target item, during the memory delay period, suggesting a serial dependence mechanism cannot fully account for the findings.

Dubé et al. (2014) reported similar findings: the task design was nearly identical to the task used in Huang and Sekuler (2010) except that responses were not made on a continuous scale. Instead, participants were presented with a probe stimulus and completed a two-alternative forced choice same/different discrimination task. Crucially, the similarity of the probe stimulus to the item in WM was manipulated across trials while the difference between the two memory items was kept constant. This design made it possible to compare the proportion of trials where participants incorrectly responded ‘same’ as a function of stimulus-probe similarity. Responses were not modulated by the target value alone, but instead by

both the target and distracter value. Incorrect 'same' responses where increased when the probe value was close to the average of the two WM items, relative to control conditions where only one WM item was presented.

Based on these findings the authors proposed that the non-target bias is the behavioural equivalent of neural averaging: when two stimuli fall within the same RFs, activation patterns reflect the average of the neural response that would be evoked by each stimulus individually. Subjects compute a weighted average of the target and distracter item, and the non-target effect on WM recall reflects an influence of the distracter on the internal target representation[3]. Such perceptual and behavioural averaging has been demonstrated repeatedly (e.g. Albrecht & Scholl, 2010; Brady & Alvarez, 2010; Parkes, Lund, Angelucci, Solomon, & Morgan, 2001; Zoccolan, Cox, & DiCarlo, 2005) but these sets of results additionally suggest that even task-irrelevant distracter stimuli, presented temporally separated from the target item, might be integrated with the target representation.

In both of these studies, effects were substantially reduced when predictive cues were presented before the memory items, suggesting that selective attention can protect against distracter interference. Huang and Sekuler (2010) proposed that selective attention to one of the two items biases the perceptual average in favour of the attended item (either by decreasing the distracter-weight or by increasing the target weight), reducing the distracter influence. Dubé et al. (2014) put forward a slightly different account: there is a greater reliance on perceptual averages under

[3] It is difficult to differentiate between an influence of distracter information on the target representations on a single trial level, or the *appearance* of such an effect in an aggregate statistic. This is a problem I will discuss later.

conditions of uncertainty, such as when item representations are weak. Selective attention during encoding reduces uncertainty about WM contents, which reduces observers' reliance on the perceptual average.

This account of interference effects on WM performance is substantially different to the feature-specific memory stores account (Magnussen, 2000). According to the perceptual averaging account, distracter information alters the target representation and shifts the stored information toward the distracter value. Interference effects emerge because information is changed, not because it is lost. There are very clear predictions that follow from the perceptual averaging account: biasing effects should be stronger the more dissimilar a distracter stimulus is and any influence should always be in the direction of the distracter value. Furthermore, the perceptual averaging account proposes that the relative weights of the target and distracter value (or the reliance on perceptual averages) can be changed by selective attention. This follows that ignored and/or task-irrelevant distracters should exert a weaker influence than attended distracter stimuli.

1.3.2 Difficulties in estimating WM contents from behaviour

Subtle effects on target representations or WM encoding are difficult to detect using binary measures of general performance such as recall accuracy. By using continuous recall tasks, where observers adjust a probe item to match the remembered target value rather than making a binary judgment, it is possible to collect WM recall error distributions and measure the quality of memory recall. For example, in an orientation memory task this approach would allow to measure whether an observer's response was

5°, 25° or 90° away from the target orientation, as opposed to simply classifying a response as correct or incorrect. Another way to uncover subtle shifts in target-related responses is to vary the similarity of the probe to the target item in a two-alternative forced choice same/different discrimination task. These approaches were used by Huang and Sekuler (2010) and Dubé et al. (2014), respectively.

However, a major concern with this approach is that it is not possible to rule out entirely the possibility that observers, on some small proportion of trials, accidentally reported a distracter item instead of the target item. Even if this only occurred on a small number of trials it would still lead to the appearance of a bias toward the distracter value in an aggregate statistic of target recall, even if such a bias does not exist at the single trial level in genuine target-related responses. The authors addressed these concerns and, on the basis of some additional analyses, argued that this is unlikely to be the case. However, a more convincing argument would be to use computational modelling approaches to estimate different sources of error in WM recall (Bays, Catalao, & Husain, 2009). Random guessing, variability in genuine target responses, and accidentally reporting a non-target value (so-called misbinding) all contribute to the overall noisiness of WM recall. Decomposing general performance into these independent contributions (Bays et al., 2009) allows not only for a more detailed understanding of WM performance, but also makes it possible to exclude distracter-related responses. Any claim that target representations are affected by distracter presentation needs to show unambiguously that effects are observed in target-related responses that are not contaminated by distracter-related responses.

1.3.3 Evidence from a neural model of WM

From the previous discussion it is evident that the exact mechanisms by which distracter information affects WM processes are unclear. It has not yet been demonstrated convincingly that distracter presentation changes the actual target representations in WM, but there is substantial evidence that distracter presentation might cause interference effects over and above causing information loss. The exact roles of selective attention, and similarity between target and distracter items, are also unclear.

A recent model for WM maintenance and encoding provides a useful framework (Johnson, Spencer, Luck, and Schöner, 2009). Although this model does not make any predictions about how incoming perceptual information shapes WM contents, it does provide a detailed account of how the processing and encoding of perceptual information is biased by WM contents. Therefore, the model provides useful constraints on how perceptual information could bias WM contents.

The model consists of a biologically plausible three-layer neural architecture that supports WM encoding and maintenance. The perceptual field receives afferent input, and provides input to the WM field. Both the perceptual and the WM field provide excitatory input to the inhibitory field, which in turn provides inhibitory input to both the perceptual and the WM field. Neighbouring neurons within the perceptual and WM field interact through lateral excitatory connections. When a stimulus is encoded, activation peaks emerge in the perceptual field and this activation, once it surpasses a certain threshold, is fed forward to the WM field. When the perceptual input stops, the activation peaks in the perceptual field disappear while the activation in the WM field is sustained. Activation peaks in the WM

field also continue to be passed on to the inhibitory field via excitatory connections. Therefore, the inhibitory field continues to provide input to the perceptual field. As a result, during WM maintenance in the absence of perceptual stimulation, the only input to the perceptual field is from the inhibitory field, driven by the activity peaks in the WM field. Those neurons coding for the feature values that are currently maintained in WM are in a suppressed state in the perceptual field. When new perceptual input is passed in, activation in the perceptual field will be suppressed if the input matches WM contents.

This model makes a specific prediction about how perceptual responses should be biased depending on the similarity of incoming perceptual information to the contents in WM. Interestingly, the model also predicts non-independence between different WM items: if two items are similar to one another, their corresponding activity peaks are close and will interact with one another. Each peak becomes more narrow and lower in amplitude as a result. When items are dissimilar, no such interaction will occur. Following the framework of this model, incoming perceptual information could only shape actively maintained representations in WM if it is different to the WM contents. Perceptual input that is identical in appearance will be suppressed by the existing activity peaks in the WM field, and is therefore not passed forward to the WM field where it could interact with existing representations. However, interactions with WM representations will not happen if incoming perceptual information is too different from WM contents, as the different activity peaks will be located away from one another.

Two recent studies allow us to test predictions of the Johnson et al. (2009) model. Sneve, Sreenivasan, Alnæs, Endestad, and Magnussen (2015) showed that holding a specific feature in mind influences neural responses to another, task-irrelevant stimulus. They asked participants to perform a WM task while recording brain activity using fMRI. Participants had to remember the spatial frequency of a grating and report whether this item, or a subsequent probe item, had the higher spatial frequency. During the delay interval, a task-irrelevant distracter stimulus was presented at the same spatial frequency, at ±1 cycles per degree (cpd) or at ±2 cpd. The variable of interest was neural responses elicited by the task-irrelevant distracter as a function of its similarity to the item maintained in memory. Distracter stimuli elicited significantly weaker BOLD responses in areas V3 and V4 when the distracter was ±1 cpd different to the WM item relative to conditions were there was no, or a ±2 cpd difference. Based on the findings, the authors concluded that those features that are task-irrelevant but immediately surrounding the task-relevant feature in feature space are in a suppressed state during WM maintenance, similar to the effects of feature-based attention on neural tuning curves and response gain reported in monkey (David, Hayden, Mazer, & Gallant, 2008a; Martinez-Trujillo & Treue, 2004) and human studies (Serences, Saproo, Scolari, Ho, & Muftuler, 2009). Behaviourally, the distracter stimulus had no impact on responses. However, observers were required to make a two-alternative forced choice to a probe stimulus, which may not be ideal for detecting subtle biasing effects, as described earlier.

Another study has similarly shown that electrophysiological responses to a non-target probe stimulus are influenced by WM contents

(Sreenivasan, Sambhara, & Jha, 2011). Participants were presented with a series of faces and asked to press a button to the presentation of a target face while their brain activity was recorded using EEG. Nontarget faces were varied in their similarity to the target item by changing the percentage of target facial features in non-target items (20% to 80%). The N250, an ERP marker of face-processing (Tanaka, Curran, Porterfield, & Collins, 2006), was largest to target items and decreased monotonically as a function of decreasing target – probe similarity. These findings suggest that the more similar WM contents are to new perceptual input, the stronger is the perceptual response.

While these studies report different results, they both converge to show that as items are maintained in WM, perceptual processing is biased in favour of the relevant item. This bias affects responses to subsequent perceptual input and the responses elicited by task-irrelevant stimuli are shaped by their similarity to the currently relevant item. However, both studies report different results regarding the influence of stimulus similarity on the effects of WM contents on perceptual processing. In Sneve et al. (2015) observers were instructed to ignore the distracter item, while in Sreenivasan et al. (2011) observers had to evaluate all stimuli, including nontarget items, to complete the task. These different task requirements may explain the different results and suggest that task-relevance plays an important role in the modulation of perceptual responses by WM contents.

While these findings show that WM contents can bias perceptual responses, the nature of the effects in both experiments is different from what is predicted by the Johnson et al. model. Sneve et al. (2015) reported a suppressed BOLD response in visual areas V3 and V4 to similar feature

values but not to identical or different feature values. According to the Johnson et al. (2009) model, however, perceptual responses should be suppressed to feature values identical to those maintained in WM. The results are also inconsistent with the distracter interference account put forward by Magnussen (2000). The findings reported by Sreenivasan et al. (2011) are even less in line with the model's prediction: responses were largest to identical stimuli and monotonically decreased as stimuli became more different. However, this was shown in an increase of the N250 responses which is related specifically to face-processing and has been localised to face-processing regions in the inferior temporal cortex (Kaufmann, Schweinberger, & Burton, 2009; Schweinberger, Pickering, Jentzsch, Burton, & Kaufmann, 2002). It is unlikely that this ERP component represents only pooled activity from the 'perceptual field'.

Irrespective of the specific details, the observation that memory contents can impact on perceptual responses is well established and these series of studies show that there are nuanced effects on perceptual responses related to the preferential tuning of population activity patterns in line with WM task requirements. If incoming perceptual information can similarly shape existing target representations in WM, the findings reviewed above suggest an influence of target-distracter similarity, task-relevance, and/or selective attention mechanisms. The perceptual averaging account predicts the greater the target-distracter difference the greater the influence, while the Johnson et al. (2009) model suggests that there is only a limited range in which WM representations interact with one another. Second, it is not clear whether distracter presentation leads to information loss, information change, or both.

1.3.4 Moderating effects of sensory signal strength

The specifics of both the perceptual averaging account and the Johnson et al. (2009) model further suggest that the signal strength of incoming information might shape biasing effects. If the reported biasing effects are related to changes in the target representations, one would expect effects to be modulated by the relative strength of the distracter item. For example, the signal strength of distracter and target items could influence the relative weights assigned to each item when computing the perceptual average (Huang & Sekuler, 2010). That is, the relative weights might not only be influenced by selective attention but also by bottom-up stimulus variables. More generally, the perceptual averaging account proposes that subjects rely more on the perceptual average when memory representations are weak, for example under conditions of high WM load (Dubé et al., 2014). However, bottom-up factors such as stimulus strength of target and distracter items could also shape the quality of their memory representations, influencing the extent to which one relies on perceptual averages.

Indeed, research has long shown that the relative contrast levels of stimuli can alter interactions between perceptual stimuli. Electrophysiological studies have shown that responses of neurons to a stimulus presented inside their RF are modulated by stimuli simultaneously presented outside the RF (Gilbert & Wiesel, 1990; Levitt & Lund, 1997). These contextual effects are stimulus-specific and depend on the orientation and direction of motion of the surround stimuli (Gilbert & Wiesel, 1990). Furthermore, Levitt and Lund (1997) showed that contextual effects are also modulated by the relative activity levels of neurons coding for the relevant

stimulus, which was changed by varying its contrast. The authors recorded single unit activity from monkey striate cortex and presented stimulus patches with preferred orientation and drift direction in a neuron's RF, and a stimulus patch of varying orientation and drift direction in the surrounding region outside the RF. When the surround stimulus was of preferred orientation and drift direction, responses to the centre stimulus were suppressed relative to a condition where the centre stimulus was presented alone. When the surround stimulus differed in orientation and drift direction, response suppression disappeared. This is the classical finding of stimulus-selective contextual effects. Next, the authors changed the contrast of the centre stimulus from an optimal, high contrast level to low contrast to influence the activity levels of V1 neurons. Contextual effects were altered by changes in the contrast level of the centre stimulus. In some cases, changing the contrast affected the stimulus-selectivity of the effect, i.e. response suppression now occurred at any surround stimulus configuration. In other cases, contrast changes did not affect stimulus selectivity but the direction of the effect: identical surround stimuli could be either facilitatory or suppressive depending on the contrast of the centre stimulus. These findings suggest that the exact details of how information is integrated depend on the relative contrasts of stimuli, that is, the relative activity levels a stimulus elicits.

The complexity of these findings also shows that varying signal strength does more than changing effects quantitatively, i.e. the influence of a low-contrast distracter signal might not simply be weaker, or at some point abolished, compared to high-contrast distracters. Instead, distracters that elicit only a weak response can lead to a qualitatively different effect

compared to distracters that elicit a strong response. Similar findings have been reported in the perceptual priming literature: prime stimuli presented below a level of conscious awareness can prime behaviour (Kiefer et al., 2011; Kouider & Dehaene, 2007). In some cases subliminal stimuli prime the same pattern of responses as supraliminal prime stimuli do, yet in other cases subliminal stimuli elicit the opposite pattern of responses compared to supraliminal prime stimuli (Eimer & Schlaghecken, 1998; 2002).

If the nature and direction of distracter effects on WM performance is changed by the relative signal strength, this would modify the perceptual averaging account (Huang & Sekuler, 2010). Although the authors do not consider subliminal information (or signal strength) specifically, the model does predict that distracting information is averaged with target information and any bias effects are in the direction of the distracter item. The perceptual averaging account would struggle to accommodate observations of low-signal or subliminal stimuli having qualitatively different effects on target representation. Studies on subliminal priming have also revealed a key role of task-relevance: subliminal information only influences behaviour when it is task-relevant but not when it was task-irrelevant (Ansorge & Neumann, 2005; Gayet, Van der Stigchel, & Paffen, 2013b). These findings suggest that distracter effects on WM contents might be modulated by the distracters' relative signal strength as well as their relevance to the task.

In chapter 2 of this thesis I ask how perceptual information biases the processing and maintenance of information WM. I ask how exposure to a distracter stimulus before or after target encoding affects WM representations, specifically whether target-related information is degraded, changed, or both. I use a mixture-model approach (Bays et al., 2009) to

exclude responses where observers accidentally report a distracter item to ensure any changes in response profiles with distracter presentation are related to effects on target-related responses, and results are not contaminated by distracter-related responses. Furthermore, I vary the relative signal strength, and task-relevance, of the distracter item to assess if, and how, such changes modulate the distracter's impact on the target representation.

1.4 Interactions within WM

From the previous literature review it is clear that many types of interactions exist between perceptual input and information stored in WM. However, other processes also shape WM. For one, WM interacts with top-down attentional mechanisms. Second, representations in memory may not be independent but interact with one another. I will consider these two topics in turn.

1.4.1 Retrospective orienting of attention in WM

A growing body of research shows that memory representations can be modulated by internal attentional selection. Even when information is no longer present in the environment, attention can be oriented retrospectively in WM and modulate responses in WM tasks: an effect termed the retrocueing benefit (Griffin & Nobre, 2003; Landman et al., 2003). Crucially, such retrocues are presented long after the lifespan of 'iconic' memory (Neisser, 1967; Sperling, 1960) and affect performance when information has already been moved into a stable WM representation. In typical retrocueing studies, observers are asked to remember an array of items and after a delay period of 500 ms (or longer) a cue indicates one item, or a

spatial location, that is subsequently probed via change detection. Performance is substantially better for validly cued items compared to neutral or invalid cueing conditions.

The exact mechanisms underlying the retrocueing benefit are still being debated. According to one account, the retrocueing benefit reflects reallocation of attention within WM, protecting cued items from decay and interference (Makovski & Jiang, 2007; Makovski, Sussman, & Jiang, 2008; Matsukura, Luck, & Vecera, 2007; Pertzov, Bays, Joseph, & Husain, 2013). It has also been proposed that retrocues decrease interference by allowing observers to remove task-irrelevant items from WM (Kuo, Stokes, & Nobre, 2012). A third account places the influence of retrocues not at the maintenance stage, but proposed that valid retrocues allow for efficient selection of items in WM (Astle, Summerfield, Griffin, & Nobre, 2012; Nobre, Griffin, & Rao, 2007).

Many studies have examined the mechanisms supporting the retrocueing benefit by fitting computational models to WM performance, and assessing which source of WM error is reduced following valid retrocues. Do valid retrocues improve the representational quality of WM items, the likelihood of successfully remembering the target, do they reduce erroneous nontarget responses (usually referred to as misbinding), or is it a combination of all of these? The results from these studies are mixed. Some studies report that retrocueing primarily increased the probability of accurate memory recall but do not influence the quality at which items are represented in WM (Murray, Nobre, Clark, Cravo, & Stokes, 2013; Myers, Walther, Wallis, Stokes, & Nobre, 2015; Souza, Rerko, Lin, & Oberauer, 2014). Others report that retrocueing affects both the probability of correctly

remembering a target item as well as the representational quality of successfully remembered target items (Gunseli, van Moorselaar, Meeter, & Olivers, 2015; Makovski & Pertzov, 2015; van Moorselaar, Gunseli, Theeuwes, & Olivers, 2015; Wallis, Stokes, Cousijn, Woolrich, & Nobre, 2015; Williams et al., 2013). Only some of these studies included a third parameter for erroneous nontarget responses, but of those that included this parameter all report effects of retrocueing on misbinding (Gunseli et al., 2015; Makovski & Pertzov, 2015; Souza et al., 2014; Souza, Rerko, & Oberauer, 2016; Wallis et al., 2015). Interestingly, a recent report by van den Berg, Awh, and Ma (2014) emphasized the importance of including a misbinding parameter in models of WM. The authors completed a factorial comparison of 32 different models of WM using data sets from different laboratories. Briefly, they found that including a misbinding parameter to WM models significantly improves model fit. Thus, to date results are mixed and it is unclear under what circumstances retrocues affect the probability of accurate target recall, the representational quality of successfully remembered target items, or erroneous nontarget reports.

Of course, there are many differences in task design between the different studies (e.g. did the task require continuous or binary responses? Was set size manipulated? Were there dual-task requirements?) and at least some of the different findings can probably be explained by changes between experimental paradigms. However, for the purpose of this thesis I would like to focus on another possibility: inherent variability in WM representations may lead to variability in observed retrocueing benefits. Indeed, the main conclusion from van den Berg et al. (2014)'s factorial model comparison was that a key aspect of WM performance – how well

each item is remembered – is variable across items and trials even for a given set size. Assuming that retrocues act on internal representations in WM, the retrocueing benefit may flexibly change with changes in the representational quality of WM items. Perhaps, some of the ambiguity regarding the functional mechanism of the retrocueing benefit is related to a flexibility of mechanisms underlying the retrocueing benefit. While there is clear evidence for variability in the quality of WM representations (Fougnie, Suchow, & Alvarez, 2012a; van den Berg et al., 2014; van den Berg, Shin, Chou, George, & Ma, 2012) the sources of the this variability are unclear. Likely candidates include fluctuations in attention, changes in task strategy and motivation across trials, and changes in stimulus properties (e.g. behavioural orientation judgments for cardinal orientations outperforms memory for oblique orientations; Girshick, Landy, amd Simoncelli, 2011). An additional bottom-up factor that would affect the quality of WM representations is stimulus contrast. If the variability in the effects of retrocues is related to variability in the representational quality of items in WM, manipulating stimulus contrast during encoding should lead to modulation of the retrocueing benefit.

All of the studies reviewed earlier used high contrast stimuli. Recently, however, a study reported effects of retrocues on a low contrast stimulus presented at threshold (Thibault, van den Berg, Cavanagh, & Sergent, 2016). Participants were asked to remember on orientation stimulus presented at threshold, and valid or invalid cues were presented before or after stimulus presentation (from -100 ms to +400 ms). After a delay period responses were made on a continuous scale. Both pre- and retrocues reduced guess rate but did not affect the quality of WM

representations. This effect was least pronounced for cues presented 400 ms post stimulus followed by cues presented 100 ms post stimulus, followed by cues presented 100 ms before stimulus onset. Importantly, most retrocueing studies had a much longer delay between presentation of the stimulus display and presentation of the retrocue. It is not clear whether the observed effect would persist beyond 400 ms. Furthermore, as only one item was presented it was not possible to assess effects of retrocueing on misbinding rates. Finally, this study only used one contrast level so it is not possible to systematically assess the impact of varying contrast levels on the retrocueing benefit.

1.4.2 Interactions of retrospective attention with signal strength

Psychophysical and electrophysiological studies have shown that effects of attentional selection on behaviour and neural firing rates vary with the relative signal and noise levels of the attended stimuli (Dosher & Lu, 2000; Lu & Dosher, 1998; 2008). For example, Lu and Dosher studied the mechanisms of selective attention by varying the level of external noise in stimuli, and then contrasted effects of attentional selection on low-noise versus high-noise stimuli. Through this approach the Perceptual Template Model (Lu & Dosher, 1998) was developed which proposes different mechanisms of attentional selection depending on the level of external noise: signal enhancement occurs during conditions of low external noise, distracter (or external noise) suppression during high levels of external noise, and contrast-gain control (modulating levels of internal noise) operates through all ranges of external noise. This topic of study is far too large to review here (see Lu & Dosher, 2008, for a review), but the key

conclusion is that there is a qualitative change of the effects of attentional selection with changes in the levels of external noise.

Others have contrasted effects of attention on neural firing rates in response to stimuli presented at different levels of contrast (e.g. Martínez-Trujillo & Treue, 2002; Reynolds, Pasternak, & Desimone, 2000). These studies have shown that attentional effects are strongly modulated by the contrast level of the attended stimulus (see Reynolds and Heeger, 2009, for a review).

Together, these findings show that the mechanisms underlying attentional benefits in perception vary with changes in the stimulus properties (external noise or contrast levels). However, none of these studies have been done in the WM domain. The effects of relative levels of signal and noise during stimulus encoding on the retrocueing benefit are unknown. In fact, the role of signal strength during encoding has been mostly overlooked in the WM literature in general, and it is unclear how this affects the representational quality of items in WM as well as other WM processes.

Some WM models have incorporated effects of internal noise levels on WM representations. Variable-precision models, for example, propose that intrinsic noise fluctuations impact on the stability of WM representations (Bays, 2014; Fougnie, Suchow, & Alvarez, 2012b; Matthey, Bays, & Dayan, 2015; van den Berg et al., 2012). Under this view, WM capacity, and quality, is determined both by the allocation of a limited mental resource to WM items as well as the stability of the information stored in memory. As a result, WM recall varies from trial-to-trial and is item-dependent. Arguably the quality of a WM representation is not just modulated by internal noise

levels immediately prior to encoding, but also by the strength of the sensory signal during encoding (as already argued above). Indeed, models of perceptual decision-making explicitly model external noise levels and assume that internal representations are shaped by both internal and external noise (e.g. Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Gold & Shadlen, 2007; Green & Swets, 1974; Macmillan & Creelman, 2005).

Other resource-based models of WM (Bays & Husain, 2008) propose that the representational quality of each item in WM depends exclusively on the amount of a limited mental resource allocated to each item, but is otherwise fixed. While there is no limit on the total number of items that can be stored in memory, the limited mental resource has to be divided over to-be-remembered items. The higher the number of items in WM, the less resource is allocated to each item. At the neural level a reduction in resource might correspond to a reduction in the gain of neural activity, for example, and the relative amplitude of the signal corresponding to a WM item against noise is reduced (Ma, Husain, & Bays, 2014). As a result, performance deteriorates as less information is available for any given item. In a scenario where a small number of items are presented at low contrast, WM capacity is not exceed but each item elicits only a weak signal. In this case, sufficient resources can be allocated to each item, but there is only a weak signal associated with each item in the first place. The relative amplitude of each signal against noise must be lower than if the same number of items were presented at a higher contrast level. How such bottom-up changes may impact on the allocation of, and demands on limited resources is unclear.

At present, it is not known whether and how attentional mechanisms underlying the retrocueing benefit differ at different levels of representational quality. If retrocueing simply supports more efficient selection of a WM representation but not the maintenance of WM items (Astle et al., 2012; Nobre et al., 2007), then retrocueing effects should not change with the robustness and quality of the target representation. However, if retrocues act during the maintenance stage, on actual stimulus representations, differences could emerge: retrocues may have more of an effect when memory representations are weak following low-contrast stimulus input. Or perhaps attentional effects only emerge once WM representations are of a certain minimal quality. Alternatively, there may not be one mechanism underlying the retrocueing benefit, but multiple mechanisms depending on the sensory quality of the information encoded in WM.

Indeed, a recent paper by Gunseli et al. (2015) showed that observers flexibly shifted retrocueing strategies along with task demands. Specifically, the retrocueing benefit seemed to rely on different mechanisms depending on the reliability of cues. Observers performed a continuous-recall task for four bars with valid, neutral or invalid retrocues. Crucially, in high-reliability blocks cues were 80% valid whereas in low-reliability blocks cues were 50% valid. Participants were explicitly informed about cue reliability at the beginning of blocks. In both 50% valid and 80% validity blocks recall precision and the probability of recalling the target item was higher for valid relative to neutral cues. However, a cueing cost, i.e. a drop in performance for invalid compared to neutral cues, was only observed when retrocues were reliable in 80% validity blocks. Misbinding was affected by retrocueing only in 80% validity blocks. Similarly, Makovski and

Pertzov (2015) found that effects of retrocues on misbinding rates were more pronounced in the presence of interference (i.e. in dual-task conditions). These findings suggest that some of the discrepancies in the retrocueing literature are due to the flexibility of mechanisms supporting the retrocueing benefit: depending on the task details certain performance patterns emerge because retrocues are used differently to improve various aspects of WM performance.

1.4.3 Rendering subliminal representations conscious

Theories of conscious awareness have proposed that bottom-up stimulus signal strength and top-down attention mechanisms interact to determine our conscious experience of the world (Baars, 1988; Baars & Franklin, 2003; James, 1890; Posner, 1994). Specifically, according to Baars' Global Workspace Theory (Baars, 1988) conscious experience requires a stimulus of sufficient signal strength as well as attentional selection of that stimulus. Phenomena such as inattentional blindness, change blindness and the attentional blink (Mack & Rock, 1998; Raymond, Shapiro, & Arnell, 1992; Rensink, O'Regan, & Clark, 1997) demonstrate how stimuli with a strong (physical) signal nevertheless escape awareness. On the other hand, extremely faint stimuli presented for short periods of time can still be detected successfully if they occur at expected points in time and/or space. Neural activation elicited by a stimulus needs to be amplified through top-down attentional mechanisms for a stimulus to reach conscious awareness, and attention has a direct role in shaping and strengthening conscious representations.

This account suggests that attentional mechanisms might render weak representations reportable, when they would have otherwise escaped awareness, even when the corresponding perceptual stimulus is no longer present in the environment. As long as the neural signal remains sufficiently strong until attentional selection occurs, the stimulus representation enters awareness once attention mechanism boost the neural signal. That is, stimuli presented below a level of conscious awareness during encoding might be rendered reportable during WM maintenance by retrospective attentional selection. Thus, if retrospective attention impacts on item representations during WM maintenance periods, this might occur in one of two ways: by changing the quality, i.e. the reportability, of an already robust and accessible WM representation, or by 'boosting' subliminal activity into a robust, supraliminal memory representation[4]. There is some evidence that the latter is the case: studies have shown that the proportion of 'not seen' judgments, and the proportion of guessing, to a threshold stimulus are reduced with retrocueing (Sergent, Wyart, Babo-Rebelo, Cohen, Naccache, & Tallon-Baudry, 2013; Thibault et al., 2016), although retrocues were only presented up to 400 ms after stimulus presentation in these studies – a delay that is much shorter than typical retrocueing studies.

[4] There are some accounts that propose WM processes do not depend on conscious awareness, and that WM can operate on subliminal representations (Hassin, Bargh, Engell, & McCulloch, 2009; Soto, Mäntylä, & Silvanto, 2011). However, here I adopt the theoretical position that WM representations and processes are, by definition, conscious (see also Stein, Kaiser, and Hesselmann, 2016, for a more detailed discussion)

1.4.4 Interactions between items in WM

Most influential models of WM assume each item in WM is independent and items do not interact with one another (Bays et al., 2009; Luck & Vogel, 1997; Zhang & Luck, 2008). However, there is now ample evidence that items in WM tend to be influenced by other items in WM (Brady & Alvarez, 2011; Brady, Konkle, & Alvarez, 2011; Matthey et al., 2015). For example, Alvarez and colleagues have shown repeatedly that ensemble statistics of a stimulus array (e.g. the spatial layout of the stimulus array or the mean size, orientation, or colour of items) biases observers' memory for individual items, suggesting that there are various representational levels within WM which interact with one another (see Brady et al., 2011, for a review).

For the purposes of the current thesis I will focus on a second example: so-called misbinding when observers mistakenly report a non-target item instead of the target item (as already discussed earlier). Although our understanding of the conditions under which misbinding occurs is limited (Matthey et al., 2015), it is a robust phenomenon that has been reported widely (Bays et al., 2009; Bays, Gorgoraptis, Wee, Marshall, & Husain, 2011; Fougnie, Asplund, & Marois, 2010; Gorgoraptis, Catalao, Bays, & Husain, 2011). According to a model proposed by (Matthey et al., 2015), misbinding errors are a consequence of representational and noise limitations in WM storage. One neural population, as opposed to multiple distinct neural populations, represents different items in WM and the representations for distinct WM items overlap. Target items, therefore, have to be extracted from overlapping representations of all items in WM. The presence of internal neural noise interferes with optimal decoding of

population activity and leads to misbinding errors (and also to increased variability in recall for the target value).

This model provides the first theoretical account of misbinding errors. However, it does not specify how neural effects of attention might modulate these processes. A study by Serences et al. (2009) used fMRI and voxel-based tuning functions to demonstrate that attention changes tuning curves: responses to attended features are enhanced whereas responses to similar but unattended features are reduced. The authors argued that this pattern reflects an increase in the signal to noise ratio of those neurons coding for the attended feature, which leads to a more robust representation. Retrospective attention might have similar effects on attended WM representations. A narrowing of tuning curves and a more robust representation of target items might reduce interference between items as target items are more easily differentiated from other, overlapping items. This would allow to read out more easily distinct items and reduce misbinding.

This explanation focuses on how attention affects the interaction between items in WM. However, as established earlier there are other sources of error in WM recall. For example, there is variability in the representation of items, so that even when a target item is correctly recalled (and not the distracter item) there is some error associated with the target representation. Attentional effects on neural tuning curves might also improve the quality of target representations, improving performance by modulating target representations rather than interactions between different representations. It is not clear under what circumstances attention might target individual representations versus interactions between items in WM.

One possibility is that this depends on the overlap or confusability of items in WM on any given trial. Items that are more similar to one another, or items that have weak, low-signal representations, might be more confusable with other items in WM. Under such circumstances retrospective attention might target interfering interactions between items.

In Chapter 4, I ask how representations in WM, and WM performance, benefit from retrospective attention and how this varies with changes in the signal strength of the to-be-encoded information. I also assess the functional role of attentional mechanisms in determining what information is experienced consciously in WM. In addition, I use the same mixture-model approach as in Chapter 2 to ask if the mechanisms underlying the retrocueing benefit flexibly shift with signal strength of the processed information, and to describe what aspects of WM performance change with variations in signal strength during encoding.

1.5 Aims of this thesis

In this thesis, I describe the flexible role of attention in information processing, and further explore ways in which information processing can be biased to shape perception and working memory. I establish a more nuanced and fuller description of the characteristics of biases in information processing. In a first set of experiments, using MEG and fMRI, I compare the neural mechanisms supporting feature-based and spatial attention. In a second set of experiments, I examine how incoming sensory information presented supraliminally or subliminally shapes WM contents. I contrast the effects on WM contents with effects on perceptual representations. In a final set of experiments, I ask how retrospective attention interacts with WM

representations to shape WM processes, including WM recall and inter-item interactions within WM. This work demonstrates that top-down mechanisms and top-down templates can flexibly modulate information processing, while these templates and mechanisms can themselves be biased and modulated throughout many stages of information processing.

2 Neural Mechanisms of Feature-based and Spatial Attention

2.1 Chapter Abstract

Selective attention can flexibly operate on the basis of feature or spatial information. Both feature-based and spatial attention are well researched in isolation, but less research has directly compared these two kinds of attention. In this chapter, I assess the extent to which neural mechanisms supporting spatial and feature-based attention overlap. I compare the two kinds of attention directly in the same experiment while recording brain activity using fMRI and MEG. At the behavioural level, feature-based and spatial attention improves performance independently and I find no evidence of interactive effects on performance. The fMRI results reveal a common fronto-parietal network for both kinds of attention, with only some regions preferentially engaged by selective spatial attention. MEG results revealed that both feature-based and spatial attention modulate alpha power during the cue-target interval, but feature-based attention modulates alpha power in a relatively more global fashion. Target-induced gamma lateralisation differentiated between feature-based and spatial attention. Enhanced lateralised gamma responses to targets occurred only after feature cues. These results suggest that both spatial and feature-based attention suppress alpha power in preparation for stimulus presentation, and top-down feature templates additionally enhance gamma responses to bottom-up stimulus input.

2.2 Introduction

Selective attention is a flexible set of mechanisms that can bias the neural processing of various aspects of incoming information. In this chapter, I will focus on contrasting and comparing spatial attention and feature-based attention, and specifically the role the low- and high-frequency oscillatory activity may play in each. Research using neurophysiological and neuroimaging methods suggests a similar network of brain areas underlying feature-based and spatial attention (Egner et al., 2008; Giesbrecht et al., 2003; Liu, Slotnick, Serences, & Yantis, 2003). This distributed pattern of brain regions has been labelled the frontoparietal control network (Corbetta et al., 1993; Corbetta, Kincade, Ollinger, McAvoy, & Shulman, 2000; Hopfinger, Buonocore, & Mangun, 2000; Kastner, De Weerd, Desimone, & Ungerleider, 1998; Mesulam, 1990; Nobre et al., 1997; Posner, Walker, Friedrich, & Rafal, 1984; Shulman, d'Avossa, Tansy, & Corbetta, 2002; Yantis et al., 2002). Recently, studies have revealed some differences in the engagement of the frontoparietal control network for feature-based and spatial attention. For example, MRI studies using multivariate analyses suggest there may be distinct subpopulations with differential tuning for feature-based versus spatial shifts of attention within the common network (Greenberg, Esterman, Wilson, Serences, & Yantis, 2010).

Some components of the network also appear to be specialized for spatial attention (Schenkluhn, Ruff, Heinen, & Chambers, 2008) and are more active following spatial cues compared to feature cues (Giesbrecht et al., 2003). Similarly, results from a combined fMRI and MEG study with

humans, and a multi-unit recording study with monkey, suggest there are regions in PFC specialised for top-down control of non-spatial (object- or feature-based) compared to spatial attention (Baldauf & Desimone, 2014; Bichot, Heard, DeGennaro, & Desimone, 2015). Despite these differences, the evidence suggests the frontoparietal control network acts as a general executive control network that supports attentional selection across a variety of domains (Egner et al., 2008; Eimer & Velzen, 2002; Giesbrecht et al., 2003; Greenberg et al., 2010; Nobre, 2004).

Research has also shown that the effects of feature-based and spatial attention on behaviour and neural activity share many similarities. At the behavioural level, selective attention leads to increased performance both at attended spatial locations or for attended stimulus features: stimuli are detected and discriminated more quickly and more accurately (Baldassi & Verghese, 2005; Carrasco, Ling, & Read, 2004; Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1990; Heinze et al., 1994; Kingstone, 2007; D. K. Lee, Koch, & Braun, 1997; Posner, 1980; Van Voorhis & Hillyard, 1977; Yeshurun & Carrasco, 1998). At the neural level, both types of attentional selection affect neural gain, neural sensitivity and neural synchronization (Fries, Reynolds, Rorie, & Desimone, 2001a; J. Lee & Maunsell, 2010; Reynolds et al., 2000). Spatial attention increases baseline firing rate and neural gain, but has little effect on neural tuning curves (Cohen & Maunsell, 2011; David, Hayden, Mazer, & Gallant, 2008b; Luck, Chelazzi, Hillyard, & Desimone, 1997; McAdams & Maunsell, 1999; Reynolds et al., 2000; Williford & Maunsell, 2006). Feature-based attention also increases baseline firing rates and neural gain (Cohen & Maunsell,

2011; David, Hayden, Mazer, & Gallant, 2008b; Martinez-Trujillo & Treue, 2004; McAdams & Maunsell, 2000). Additionally, changes in feature selectivity by modulating neural tuning curves have also been reported (David, Hayden, Mazer, & Gallant, 2008b; Haenny & Schiller, 1988; Motter, 1994; Ogawa & Komatsu, 2004).

In addition to the modulation of spiking activity in neurons, selective attention has also been proposed to modulate the integration and routing of neuronal activity across brain regions through changes in neural synchronisation. Indeed, neural synchronisation of oscillatory activity has been proposed to contribute to the physiological mechanism for directed attention (Jensen et al., 2007; Klimesch, 2012; Womelsdorf & Fries, 2007). Since the first seminal demonstration that selective attention increases neuronal synchronisation associated with high-frequency gamma-band activity and decreases low-frequency activity (Fries et al., 2001), the role of oscillatory synchronization in high frequency (gamma, > 30 Hz) and low frequency (alpha, 8 – 14 Hz) bands in selective attention tasks has remained an active area of research. There are, broadly speaking, two ways in which selective attention can improve performance: for one, performance can be improved via inhibition of task-irrelevant information and neural populations that code for unattended features or locations (Slotnick, Schwarzbach, & Yantis, 2003; Smith, Singh, & Greenlee, 2000; Vanduffel, Tootell, & Orban, 2000). Additionally, performance improvement can be achieved by enhancing the processing of task-relevant information in neurons processing the attended stimulus (Foxe, Simpson, Ahlfors, & Saron, 2005; Kastner, Pinsk, De Weerd, Desimone, & Ungerleider, 1999;

Müller, Teder-Sälejärvi, & Hillyard, 1998). Low frequency oscillation modulations have been linked to both inhibition and excitation functions, and modulations of high frequency oscillations have been linked to excitation functions (Fries, Reynolds, Rorie, & Desimone, 2015; Jensen & Mazaheri, 2010; Kelly et al., 2006; Klimesch, 2012; Klimesch et al., 2007).

Modulations in alpha power are a hallmark of spatial attention studies. Alpha power modulations are involved in anticipatory biasing of attention and reflect the allocation of attentional resources. Modulations typically peak during the anticipatory delay period just before stimulus onset. In spatial cueing tasks, alpha-band power is reduced contralaterally to the attended region of space during the cue-target interval relative to baseline both in monkeys (Bollimunta et al., 2011; Buffalo et al., 2011; Fries et al., 2008; Fries, Reynolds, Rorie, & Desimone, 2001a; Gregoriou, Gotts, Zhou, & Desimone, 2009b) and humans (Bauer et al., 2006; Bauer, Kennett, & Driver, 2012; Gould, Rushworth, & Nobre, 2011; Sauseng et al., 2005; Siegel et al., 2008; Thut et al., 2006; Wyart & Tallon-Baudry, 2008b). Additionally, alpha-band power has also been shown to increase ipsilaterally to the unattended region of space relative to baseline levels (Kelly et al., 2006; Worden et al., 2000). According to one framework, a decrease in alpha power reflects a release from inhibition of those neurons coding for attended regions of space and features, and an increase in alpha reflects active suppression of those neurons coding for unattended regions of space and features (Jensen & Mazaheri, 2010; Klimesch et al., 2007).

In feature-cueing tasks similar patterns of alpha-band modulations have been observed: alpha power decreases relative to baseline in cortical

regions selective for task-relevant feature dimensions, and increases relative to baseline in neurons coding task-irrelevant feature dimensions (Müller & Keil, 2004; Snyder & Foxe, 2010). In both studies, feature-based effects were observed bilaterally in each hemisphere suggesting that for feature-based attention there is a global modulation of alpha power that extends across the cortex to target spatially dispersed neurons.

Gamma-band activity, on the other hand, has been implicated in the active processing of visual stimuli. Studies typically report a relative increase in gamma power in cortical regions processing attended locations or features. For example, electrophysiological recordings in monkey areas V1, V2 and V4 have shown that spatial attention increases gamma-band synchronization contralateral to the attended hemifield primarily following stimulus presentation and to some extent also during the pre-stimulus delay period (Buffalo et al., 2011; Fries et al., 2008; Fries, Reynolds, Rorie, & Desimone, 2001a; Womelsdorf et al., 2006). Others have reported similar effects of spatial attention in humans using EEG and MEG recordings (Bauer et al., 2006; Gruber et al., 1999; Koelewijn, Rich, Muthukumaraswamy, & Singh, 2013; Landau, Esterman, Robertson, Bentin, & Prinzmetal, 2007; Siegel et al., 2008; Wyart & Tallon-Baudry, 2008b).

Feature-based attention has also been shown to increase gamma-band power in monkey area V4 in those neurons that code the attended feature, irrespective of the locus of spatial attention (Bichot et al., 2005). Similarly, EEG studies with humans have shown bilateral increases in gamma power for attended relative to non-attended features (Müller & Keil, 2004).

The interpretation of alpha modulation as a mechanism that inhibits task-irrelevant information is relatively intuitive. The available evidence fits this proposal, but the proposal is not undisputed (e.g. Bauer, Stenner, Friston, & Dolan, 2014b). The interpretation of gamma power modulations is less clear. The role of gamma-band modulations is not limited to selective attention and enhancing neural processing of task-relevant stimuli. Gamma-band activity has been proposed to be involved in a variety of functions other than attention, including feature binding and coherent object perception (Eckhorn et al., 1988; Singer & Gray, 1995; Tallon-Baudry et al., 1996; 1997), memory (Honkanen et al., 2014; Pesaran et al., 2002; Pulvermüller et al., 1999; Tallon-Baudry et al., 1998) and visual awareness (Wyart & Tallon-Baudry, 2008b). Considering the multitude of purported gamma-band functions, recent proposals have been made that gamma power modulations represent a fundamental functional mechanism of the brain. This mechanism allows for precise and optimal temporal coding of information relay, and as such is critically involved in a variety of higher cognitive functions of which selective attention is only one (Fries, 2009; Tallon-Baudry, 2009; Vidal et al., 2006).

While both feature-based and spatial attention are relatively well studied in isolation, only a few studies have directly compared them. White, Rolfs, and Carrasco (2015) compared the behavioural consequences of feature-based and spatial attention directly using a target-detection task with spatial and feature cues. Participants were presented with superimposed groups of red and green dots on the left and right side, each group extending over the upper and lower visual field. They were instructed to

report the location (upper or lower visual field) of a saturation change in one of the groups of dots. Different cueing conditions were employed to manipulate feature-based and spatial attention separately: valid, neutral or invalid feature (red or green dots) or spatial (left or right side) pre-cues were presented before the target event. Stimulus competition during the saturation change was also varied: in the first experiment four simultaneous saturation changes occurred with three serving as distractors, whereas in the second experiment only one saturation change occurred. Post-cues presented after the target event, before the response was made, ensured that uncertainties regarding the target event's side or colour were equated across conditions. Colour and spatial cueing effects on D prime and RTs were independent and additive in Experiment 2 when there was no stimulus competition. However, in Experiment 1, with stimulus competition, colour and spatial cueing effects interacted and largest cueing effects occurred when both spatial and colour cues were valid. The authors then developed a novel computational model (the Independent Systems with Competition Model, ISC) to simulate the behavioural pattern of performance. The ISC assumes independent, additive feature-based and spatial attentional modulations along with a normalization process that implements competition between multiple stimulus representations. Specifically, as part of the normalisation each stimulus response gets raised to an exponent before being divided by the sum of all stimulus responses. This happens after attentional amplification of stimulus responses and consequently exponentiation exaggerates any existing differences between stimulus responses in a non-linear manner. The model was able to produce the

same pattern of results: interactive effects under conditions of stimulus competition, and additive effects otherwise. From this, the authors argued that feature-based and spatial attention operate independently and that independent attentional effects can develop into interactive effects when competition introduces nonlinearities.

Giesbrecht et al. (2003) tested the generalizability of the frontoparietal control network in attentional control by directly contrasting spatial and feature-based attention within the same experiment. Participants performed a cued attention task and oriented their attention to either a spatial location (left or right) or a colour (blue or yellow) while brain activity was recorded using event-related fMRI. In spatial-cue trials, two probe stimuli of the same colour were presented bilaterally, while in feature-cue trials the two probe stimuli were shown overlapping at the centre of the screen, one in blue and one in yellow. The task was to indicate the orientation of the probe stimulus of the cued dimension, i.e. the cued location or the cued colour. Highly similar, overlapping brain regions encompassing the frontoparietal control network were activated in response to feature and spatial cues, including dorsal frontal areas and superior and inferior parietal areas. Some regions were significantly more active following spatial cues compared to feature cues suggesting that at least some components of the network were preferentially engaged by attentional orienting in space. These results suggest that the frontoparietal control network is not specific to spatial attention but instead is a general executive network that is recruited for both feature-based and spatial attention, in line with the conclusions drawn from previous studies that examined feature-

based and spatial attention separately (see also Kanwisher & Wojciulik, 2000).

Cohen and Maunsell (2011) compared the effects of feature-based and spatial attention on neuronal activity by recording from monkey area V4 bilaterally. Monkeys completed a change-detection task in which attention was manipulated by cueing either the feature dimension or the location of the change. The effects of feature-based and spatial attention on local populations of neurons were comparable: firing rates to preferred stimuli were increased with attention while firing rates to non-preferred stimuli were decreased. However, the cortical extent of feature-based and spatial attention effects differed: feature-based attention was coordinated across hemispheres and targeted spatially dispersed neurons. In contrast, spatial attention targeted only spatially localized neurons and spatial attention effects across hemispheres were not coordinated. The authors concluded that feature-based and spatial attention rely on similar mechanisms but are separate processes that operate on different spatial scales.

To summarise, findings of studies comparing and contrasting spatial and feature-based attention using psychophysical, neuroimaging and electrophysiology methodologies, converge on the conclusion that feature-based and spatial attention operate largely independently but rely on similar mechanisms. Crucially, feature-based and spatial attention have to operate across different spatial scales because of their different neural representations. The observed similarities between feature-based and spatial attention led to the proposal of the ‘feature-similarity gain model’ (Treue & Martínez Trujillo, 1999) according to which the responses of

sensory neurons reflect the similarity of spatial and non-spatial feature of the current target to the sensory selectivity of the neurons. In other words, if an attended location matches neurons' receptive fields, or a feature their sensory selectivity, neuronal responses will increase. Because visual input is retinotopically organised, spatial attention can operate by modulating neighbouring groups of neurons in a retinotopic fashion (Hopfinger et al., 2000; Luck et al., 1997; Maunsell & Treue, 2006; Tootell et al., 1998; Woldorff et al., 1997). Targeted neurons are located near each other and correspond precisely to the attended location. The exact nature of the neural representation of features, on the other hand, is less clear: some features area coded topographically but this is unlikely to be the case for all feature dimensions. While our understanding of feature representations is incomplete, it is clear that features are coded in a spatially dispersed fashion throughout the cortex. Spatial attention can, therefore, rely on a local retinotopic mechanism, while feature-based attention requires a more global mechanism that can target neurons across the whole visual field.

In addition to the results discussed earlier, there is a growing body of psychophysical, neuroimaging and electrophysiological studies that support the notion that feature-based attention operates globally, improving the processing of task-relevant features throughout the entire visual field (Bichot et al., 2005; Jehee et al., 2011; Martinez-Trujillo & Treue, 2004; McAdams & Maunsell, 2000; Saenz et al., 2002; Serences & Boynton, 2007; Treue & Martínez Trujillo, 1999). Treue and Martinez-Trujillo (1999) recorded from macaque area MT and presented animal with two motion direction stimuli at two locations, one inside the RF of the neurons being recorded and one

outside the RF. Animals were trained to attend to the motion direction of the stimulus outside the RF. The ignored stimulus inside the RF always moved in the neuron's preferred direction. The attended stimulus outside the RF moved either in the preferred or in a non-preferred direction. Responses to the ignored stimulus were affected by the changes in attended motion direction: responses were increased when the same, preferred motion direction was attended relative to when a different, non-preferred motion direction was attended. McAdams and Maunsell (2000) reported similar results in macaque area V4. Responses to a distracter orientation stimulus were increased when an identical orientation stimulus was attended at a different location, relative to when a coloured patch was attended at a different location. Using a visual search paradigm, Bichot and colleagues (2005) also showed global modulations on stimulus processing by feature-based attention in monkey area V4. Animals were trained to search for a cued target item that consisted of two features (shape and colour) and neural responses were compared between trials where the search target fell in the RF but was not detected, and trials where the same stimulus fell in the RF but it was not the search target. Responses were increased when the stimulus in the RF matched the cued feature, i.e. it was the search target, relative to when it was not the search target, and this effect was independent of the locus of spatial attention. In sum, electrophysiological recordings from single neurons suggest that attention to a feature increases responses of neurons sensitive to that feature globally across the visual field independently from spatial attention, even to task-irrelevant stimuli.

Saenz, Buracas and Boynton (2002) demonstrated global effects of feature-based attention using fMRI in human participants. Participants completed an adapted version of the task used in Treue and Martinez-Trujillo (1999): participants performed a speed-discrimination task in one of two superimposed fields of upward and downwards moving dots, which were presented on one side of the visual field. Simultaneously, on the other side of the visual field, a single field of upward- or downward-moving dots was shown which observers were instructed to ignore. BOLD responses to ignored stimuli in visual area MT+ were increased when it moved in the same direction as the attended field in the target stimulus, supporting the conclusion that FBA modulates neural activity in a spatially global manner across multiple early stages of visual processing. Jehee, Brady and Tong (2011) and Serences and Boynton (2007) used fMRI and pattern classification methods to show global effects of FBA whereby effects spread to task-irrelevant locations in the visual field.

Finally, global effects of feature-based attention across the visual field also find support in behavioural and neuropsychological studies: the contingent capture effect shows that behaviourally-relevant features attract attention even when they are presented at unexpected locations (Folk et al., 2002; Serences et al., 2005). Saenz, Buracas, and Boynton (2003) showed that performance in a dual task was better when spatial attention was divided over stimuli that shared a common feature relative to when they had different features. Together, there is converging evidence from behavioural, neuroimaging, and electrophysiological studies that feature-based and

spatial attention operate across different spatial scales, but rely on similar mechanisms.

As noted, high- and low-frequency oscillations have been implicated in both types of attentional selection, however their relative contributions in each case have not been directly compared. To address this question, I asked participants to complete an attentional cueing task that manipulated feature-based and spatial attention factorially, while I recorded brain activity using MEG. This design allowed me to examine effects of feature and spatial cues separately as well as their possible interactions. I focused the analyses on the alpha-band during the cue-target interval and on the stimulus-related gamma-band. In a separate session, brain activity was recorded using fMRI on a simplified version of the task in order to compare the network of brain areas engaged in each case.

Based on the previous literature, I predicted independent effects of feature-based and spatial attention at the behavioural and neural levels. I expected to observe lateralised alpha power modulations following spatial cues, and non-lateralised effects of feature-based cues throughout the visual field. Specifically, comparing cues carrying only spatial information to combined spatial-feature cues, I predicted additional bilateral alpha power decreases for combined cues. I also expected to observe an effect of isolated feature cueing on alpha power that was independent of spatial cues whereby alpha power is decreased bilaterally following feature cues relative to cues without feature information. In the gamma-band, I predicted an increase of power to target presentation in spatial cueing trials relative to neutral cueing trials. Gamma power has also been implicated in feature-

specific memory maintenance (Honkanen et al., 2014) and to the matching of memory contents to bottom-up input (Herrmann et al., 2004). Since feature attention, unlike spatial attention, is likely to involve the maintenance of a feature template in visual working memory, I expected to observe an increase in gamma power following feature cues relative to non-feature cues during the preparatory period prior to target presentation, and/or to target presentation. Finally, I predicted comparable patterns of activation in the frontoparietal control network following feature and spatial cues.

2.3 Methods

2.3.1 Participants

All experimental protocols were reviewed and approved by the Central University Research Ethics Committee of the University of Oxford. Twenty students took part in the experiment (9 male, 18 – 35 years, all right-handed). All participants had normal or corrected-to-normal vision and were naïve to the purpose of the experiment. Participants gave informed consent before taking part, and received course credits or financial compensation (£15/hour for sessions involving MEG or MRI and £10/hour for sessions involving only behavioural testing) for taking part.

2.3.2 Task and Stimuli

In order to examine the role of brain rhythms in feature-based and spatial attention, I asked participants to perform a target-detection task on two peripheral orientation stimuli, which were preceded by a cue. Figure 2-2A provides a task schematic of the experimental task. Most participants completed three experimental sessions. An initial behavioural session was

conducted to provide training on the behavioural task. Task performance was accompanied by MEG recordings in the second session. An additional session was performed in the MRI scanner to collect structural and functional scans. Tasks were identical in the training and MEG session, but an adapted version was used in the MRI session with the aim to make the task both shorter and easier.

2.1.1 Experimental Task

Stimuli were created on MATLAB v.7.10 (MathWorks) and presented using the Psychophysics Toolbox in Matlab (Kleiner, Brainard, & Pelli, 2007) against a uniform mid-grey background on a 23-inch light-emitting-diode (LED) display with a spatial resolution of 1920x1080 and a vertical refresh rate of 120 Hz.

The task involved detecting a peripherally presented (left or right visual field) target orientation stimulus (horizontal or vertical grating) preceded by a centrally presented spatial and/or feature cue. A distracting tilted foil stimulus occurred on the other visual field. Participants were instructed to maintain central fixation during task performance.

Trials began with a 1,000 – 1,500 ms (randomly determined) presentation of a central fixation dot. This was followed by the presentation of a central cue stimulus for 200 ms. The cue stimulus consisted of a diamond shape (1.59° visual angle in width and height) around a plus-sign in its centre. To cue spatial location (left vs right) the left or right half of the diamond shape was highlighted in blue or magenta, forming a leftward or rightward pointing arrow. To cue the feature identity (horizontal vs vertical)

the vertical or horizontal line of the plus sign was highlighted in blue or magenta. The highlight colour was counterbalanced between participants, with the other colour serving as the default colour of the cue stimulus (i.e. the neutral cue was either all blue or all magenta). Participants were explicitly told that cues did not predict whether the target would be present or absent. Following a cue-target interstimulus interval (ISI) of 1,200 ms, a stimulus array was presented for 50 ms, and then backwards-masked (after an interval of 67 ms) for 283 ms. Stimulus arrays consisted of either a horizontally or vertically oriented target Gabor patch (Gaussian-vignetted sinusoidal gratings) and a ±45° tilted distracter Gabor patch (target present), or two distracter tilted Gabor patches (target absent). Gabor patch stimuli (diameter: 2° of visual angle) were presented at a contrast for which the target detection task was performed at 75% accuracy (see below for details of the Staircasing procedure) with a spatial frequency of 3 cpd. The Gaussian envelope had a space constant of 0.44°. The orientations of the target and distracter stimuli were randomly selected on each trial. A target was present on half the trials (randomly selected at the beginning of the experiment). Gabor patches were presented ±5.19° from the vertical meridian and 3° below the horizontal meridian. Backwards-mask stimuli were constructed by applying a Gaussian-vignette to the convolution of 100% contrast square-wave gratings at four orientations of 90°, 180°, 45°, and -45° angle. Stimuli and masks were presented atop 10% contrast luminance pedestals, which were present throughout the entire experiment. Participants responded (target present or absent) by making a right-handed index or middle finger button press during a response period of up to 2,000

ms. Response mappings were counterbalanced between participants. The subsequent inter-trial interval (ITI) was 1,000 ms (±500 ms). All conditions were equiprobable and randomized.

The task was adapted for the MRI session as follows: the fixation duration was shortened to 500 ms, the cue-target ISI was shortened to 600 ms, no backwards mask was presented, the response period lasted up to 1,000 ms, and the ITI was 200 ms. Furthermore, the Gabor patch stimuli were presented at 50% contrast.

2.4 Procedures

Participants completed the behavioural and MEG sessions on separate days. A subset of participants (N = 15) completed an additional MRI session on a third day. Participants underwent a calibration session before the behavioural and MEG experimental session using an adaptive psychophysical staircase procedure to estimate the threshold contrast for perceiving the Gabors. Task difficulty was adjusted for each participant by titrating the contrast of the Gabor patches for which the target detection task was performed at 75% accuracy when only neutral cues were used.

2.4.1 Behavioural training session

Participants completed a total of 288 trials of the experimental task split into three runs. In each run short rest periods were introduced every 32 trials. The first run served as a practice run to familiarise participants with the task, and stimulus contrast was set to 50%. All cues were included in this run. The second run comprised the staircase procedure starting at 50%

contrast, and was repeated if necessary. Only neutral cues were included in the staircase run. In the third run, the contrast was set to the threshold estimated from the staircase, and spatial, feature, and spatial-feature cues were included.

2.4.2 fMRI session

Magnetic-resonance images were obtained with a 3T TrioTim Siemens system (Siemens Medical Solutations, Erlangen, Germany). Functional measures sensitive to the blood oxygenation level-dependent (BOLD) contrast were obtained using single-shot echo planar T2*-weighted imaging (TE = 30 ms, TR = 3s, Flip Angle = 87°). Forty-five 3-mm slices (64 x 64 voxel matrix, with a 3 mm^2 in-plane resolution) were acquired with an interleaved ascending sequence, covering the entire cortex. Participants completed three experimental runs (with one exception where only two experimental runs were completed) each containing 129 image sets. The first four image sets were collected in the absence of any task to allow the signal intensities to saturate and were discarded from subsequent analysis. Structural magnetic-resonance images were obtained during the same experimental session using a high-resolution T1-weighted sequence (1 x 1 x 1 mm resolution, TE = 4.7ms, TR = 2.04s, Flip Angle = 8°).

In each run, participants completed 96 trials split into six blocks of 16 trials each. There were two blocks for each of the three different cueing conditions used (spatial, feature, and neutral; the spatial-feature cue condition was not included). Task blocks lasted 42 s, and were separated by 18 s of rest. Block order avoided the immediate repetition of condition, and block-order was counterbalanced between participants. In total, participants

completed 288 trials over the three runs, resulting in 96 trials for each cue type.

2.4.3 MEG session

Whole-head MEG recordings were acquired with a 306-channel Vectorview system (Elekta-Neuromag, Helsinki, Finland) at the Oxford Centre for Human Brain Activity (OHBA). The system contains 306 sensors: 102 magnetometers and 204 orthogonally oriented planar gradiometers (102 latitudinal gradiometers, 102 longitudinal gradiometers). The MEG scanner was housed in a three-layer magnetically shielded room (ImedcoAG). The data were sampled at 1000 Hz.

Additional electrodes were used to record eye movements (EOG) and heart rate (ECG) during MEG acquisition. A vertical pair of EOG electrodes was placed above and below the left eye to detect blinks. To detect lateral eye movements, a horizontal pair was placed with one electrode to the left of the left orbit, and the other electrode to the right of the right orbit. Electrodes were placed on the left and right wrist to record ECG. In addition, eye movements were recorded at 1,000 Hz with a remote infrared eye-tracker (SR research, EyeLink 1000) controlled via the Psychophysical Toolbox in Matlab v.7.10 (MathWorks). Four magnetic coils served as head-position indicator (HPI) coils. These were positioned on each mastoid bone and on each side of the forehead near the hairline. The positions of the HPI coils relative to 3 fiducial points, as well as the subject's headshape, were digitized using a Polhemus 3D tracking system (Polhemus, EastTrach 3D). HPI coils were activated at the beginning of each block to localize the participant's head with respect to the sensor array

and to monitor the subject's head position to correct for any head movements later during analysis.

During the recording session, participants sat in a reclining chair and supported their head against the back and top of the MEG cap. Participants were asked to remain as still as possible during the recording session and were continuously monitored by a video camera. Participants were given a fibre-optic button box and made responses using their right and left index finger. Stimuli were projected onto a screen placed around 120 cm in front of the subject. MEG data were recorded in 5 blocks. Participants completed a total of 640 trials of the experimental task split into five blocks of 128 trials each. There were four different conditions (feature-spatial cue, spatial cue, feature cue, neutral cue), which were presented in randomized order, resulting in a total of 160 trials per condition.

2.5 Analysis

2.5.1 Behavioural Analysis

Hypotheses regarding the effects of experimental parameters on recall accuracy and reaction times (RTs) in correct trials were tested using analysis of variance (ANOVA) and *t* tests (Bonferroni corrected where appropriate) for participants' performance in the MEG and MRI session separately. Performance in the training session was not analysed.

2.5.2 fMRI analysis

Data were processed and analysed using Statistical Parametric Mapping (SPM 8, Wellcome Department of Cognitive Neurology, London, UK) implemented in Matlab (MathWorks, Natick, MA). Statistical maps were

projected onto a cortical surface template using Caret software and the population average landmark surface (PALS, Van Essen, 2005), or, for the axial view, using MRIcron (http://www.sph.sc.edu/comd/rorden/mricron/) overlaying statistical maps on a standard template. Functional scans were corrected for differences in slice timing, and then realigned and unwarped to correct for movement artefacts. High-resolution anatomical T1 images were spatially coregistered with the realigned functional images to enable the localization of functional activations relative to the structural anatomy. The coregistered structural images were segmented and transformed into standard Montreal Neurological Institute (MNI) space using a unified segmentation approach (Ashburner and Friston, 2005). Functional images were spatially smoothed using an 8-mm Gaussian kernel to conform the data to a Gaussian model and to accommodate inter-subject anatomical variability.

For first-level analyses, task-related activity was modelled by convolving a vector of the trial onset times with a synthetic hemodynamic response and its temporal derivative. Statistical analyses were implemented in the General Linear Model using SPM 8 modelling effects of interest and other confounding effects (e.g. session effects and motion). There were 11 first-level regressors: 5 task variables (left cue, right cue, horizontal cue, vertical cue, neutral cue) and 6 motion regressors (three translations, three rotations). Group analyses were conducted using random-effects models to assess the effect of cue type. The following contrast images were taken to the second level: 1) contrasts comparing overall task related activity to baseline activity (all task variables > 0), 2) contrasts comparing spatial

preparatory attention to feature preparatory attention (left/right cues > horizontal/vertical cues), 3) contrasts comparing feature preparatory attention to spatial preparatory attention (horizontal/vertical cues > left/right cues), 4) contrasts comparing spatial preparatory attention to general cue-related activity (left/right cues > neutral cues), and 5) contrasts comparing feature preparatory attention to general cue-related activity (horizontal/vertical cues > neutral cues). I did not include a contrast comparing horizontal to vertical cues as I predicted that a univariate voxel-wise analysis would not be sensitive to differences in activation between the two conditions. All SPM{Z} values were thresholded at a voxel-level uncorrected $P \leq 0.001$ ($Z = 3.09$). The threshold for reporting a significant cluster was cluster-level $P \leq 0.05$ corrected for multiple comparisons over the entire brain volume (Poline, Holmes, Worsley, and Friston, 1997).

2.5.3 MEG analysis

Analysis of the MEG data was performed using custom-written MATLAB scripts, the in-house OHBA software library (OSL), SPM8 (Litvak et al., 2011, http://www.fil.ion.ucl.ac.uk/spm), and FieldTrip (Oostenveld et al., 2011; http://www.ru.nl/fcdonders/fieldtrip/). All time samples were corrected with respect to the refresh delay of the projector (measured on-line with a photodiode). MEG analysis was performed in three main steps: 1) preprocessing of the data to remove artefacts present in the raw data, 2) sensor-level analysis consisting of the time-frequency decomposition of the MEG data to examine the temporo-spectral dynamics of alpha-band and gamma-band modulations, and 3) sensor-level analysis of event-related fields (ERFs) time-locked to the stimulus display. The epochs of interest for

analyses were the period between cue and probe onset, and the period immediately following probe onset.

Preprocessing

The raw MEG data were visually inspected to remove any channels with excessive noise. The data were then de-noised using Maxfilter signal-space separation (Taulu, Simola, & Kajola, 2005), and compensated for changes in head position within-session using the MaxMove software, both implemented in MaxFilter version 2.2 (Elekta Neuromag). The Maxfilter software works by mathematically transforming the data to a set of virtual sensors, which is made possible by location information provided by the MEG sensor array and each person's continuous head position signal recorded from the HPI coils. Transforming the data into virtual sensors gives a standard reference frame and allows for data to be combined across recordings. The data were subsequently down-sampled to 250 Hz. After applying a high-pass filter at 1 Hz, the data were then epoched with respect to each cue onset (from -2 s to + 3 s). This time window encompassed both cue-probe interval activity (from 0 s onward) and probe-related activity (from 1.4 s onward). The resulting epochs were manually inspected for artefacts. Systematic artefacts (including eyeblinks and heart beats) were identified and regressed out of the data using the following procedure. Independent component analysis (ICA) was used to decompose the sensor data for each session into 150 temporally independent components (tICs) and associated sensor topographies using FastICA (http://research.ics.aalto.fi/ica/fastica). Artifact components were manually identified by the combined inspection of the spatial topography, time course, kurtosis of the time course and

frequency spectrum for all components. Artefacts were then rejected by subtracting them out of the data. Finally, data were imported into Fieldtrip and inspected using the semi-automatic rejection tool to discard any remaining trials with excessive variance at 8 – 14 Hz (for alpha) or 55 – 100 Hz (for gamma).

Sensor space analysis

Time-frequency analysis of the time-domain sensor-space signal was performed using a Fast Fourier transformation with separate parameters for low (3 – 30 Hz) and high (40 – 120 Hz) frequencies. The interval -1 s to + 2.5 s centred on cue onset was used for analysis. For low frequencies, Fourier-transformed data were multiplied with a Hanning taper in 1-Hz steps using a fixed 300-ms sliding time window moving in steps of 60 ms. Power of high-frequency oscillations was estimated using a multi-taper frequency transformation with a set of discrete prolate spheroidal tapers (Percival and Walden, 1993) in 2-Hz steps using a 300-ms sliding time window moving in steps of 60 ms. I then averaged the resulting power spectra over trials within each condition of interest. Magnetometers were discarded and the power time-series in the gradiometer pairs were combined (Cartesian sum) resulting in a 102 sensor combined planar gradiometer map of sensor space power.

For the low-frequency analysis, for each participant I contrasted the power spectra in the left minus right spatial cue conditions to examine alpha lateralisation following spatial cues. Additionally, to examine effects of feature information on cue-induced alpha lateralisation I calculated a lateralisation index for spatial cues with and without feature information:

[(Contralateral responses - ipsilateral responses) / (contralateral responses + ipsilateral responses)]*100, i.e. responses were collapsed over left and right channels. I then compared lateralisation indices for spatial cues without feature information to spatial cues with feature information.

Second, I examined non-lateralised alpha power. I contrasted power spectra in feature (spatial-feature and pure feature cues) minus non-feature (pure spatial and neutral cues) cueing conditions to examine the main effect of feature cueing on alpha power. Then, to examine feature modulation of the spatial cueing effect, I contrasted the power spectra in the spatial cue with feature minus spatial cue without feature conditions, separately for left and right spatial cues, e.g. left /right spatial cues with feature minus left/right spatial cues without feature.

For the high-frequency analysis, to reduce dimensionality of the data, I based my analysis on averages over central-posterior channels (see Figure 2-1A). I examined effects of spatial and feature cueing on target-induced gamma lateralisation. I calculated lateralisation indices: [(Contralateral responses - ipsilateral responses) / (contralateral responses + ipsilateral responses)]*100 for spatial (spatial-feature and pure spatial cues), non-spatial (pure feature and neutral cues), feature (spatial-feature and pure feature cues), and non-feature (pure spatial and neutral cues) cueing conditions. Then, I compared feature to non-feature lateralisation indices and spatial to non-spatial lateralisation indices. Laterality was calculated relative to the side of target presentation, i.e. only target-present trials were included in the analysis.

At the group level, I tested for significance using paired t-tests. I used spatial cluster permutation statistics (5000 permutations, alpha-level .05) to control for multiple comparisons. Sensor-space cluster permutation statistics were computed by permuting condition labels (e.g. left and right cue conditions) using Fieldtrip's ft_freqstatistics. For low-frequency analysis, clusters were formed in space (channels) and time (0 ms to 1,400 ms from cue onset), averaging over the alpha band (8 – 14 Hz), and tested for significance against the permuted distribution. For the high-frequency analysis, clusters were formed in time (1,400 ms to 2,000 ms), averaging over the gamma band (40 – 100 Hz).

Event-related fields

Event-related fields (ERFs) are the magnetic equivalent of event-related potentials as measured by EEG. Epochs for the ERF analysis were filtered with a low-pass filter of 40 Hz, magnetometers were discarded, and epochs then averaged to provide an average waveform for contralateral and ipsilateral sides relative to the side of target presentation (collapsed over left and right side). I examined latitudinal and longitudinal planar gradiometer separately and did not combine the power time-series in the gradiometer pairs. Waveforms were compared between contralateral and ipsilateral responses evoked by targets following spatial cues (both with and without feature information). I also compared contralateral and ipsilateral target-evoked responses following neutral cues, feature cues, spatial cues, and combined spatial-feature cues. Epochs were normalised using a baseline from -400 to 0 ms. Analysis focused on a posterior channel selection over visual areas. Because of non-negligible inter-individual differences in head

size and positioning in the MEG helmet, somewhat different channels will record activity over visual areas for different subjects. Since alpha lateralisation is generally assumed to reflect modulation of activity in visual areas, I used alpha lateralisation as an index to identify channels with activity coming from visual cortex. Specifically, I selected channels above left and right visual areas by plotting alpha lateralisation (from the preceding time-frequency analysis, averaged over all types of spatial cues) for individual subjects and manually selecting the left and right channels that showed the maximal response (power suppression contralateral and increase ipsilateral). For a small number of subjects (N = 4) there was no obvious induced alpha-band modulation in the cue-probe interval, and I therefore based channel selection on the grand average alpha lateralisation (see Figure 2-1B for examples). All ERF analyses are based on contrasts that are orthogonal to the contrast used for channel selection. Specifically, channel selection was based on alpha modulation during the cue-target interval following left versus right spatial cues independent of feature information. ERF analyses examine effects of spatial cues on target-evoked responses as a function of feature cueing.

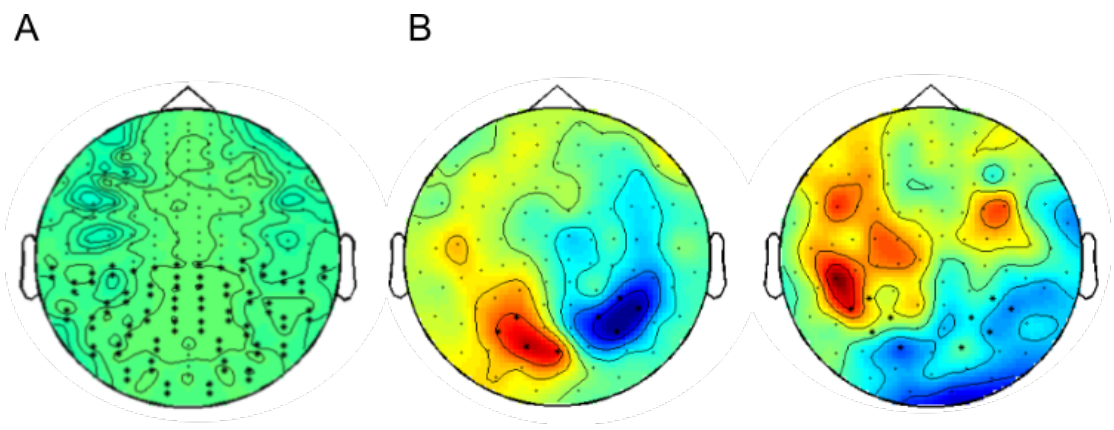

Figure 2-1. A. Central-posterior channel selection used in all gamma-band analyses. **B.** Channel selection used in ERF analyses. For most subjects selection is based on individual alpha-lateralisation during the cue-target interval (see example on left). For all small selection of participants (n = 4) channel selection was based on the grand average of cue-target alpha-lateralisation, as they showed no obvious alpha-band modulation (see example on right).

2.6 Results

2.6.1 Behavioural Results

fMRI session

Repeated-measures analyses of variance (ANOVAs) were carried out on accuracy and RTs (in ms) from correct trials with the factor of cue (spatial, feature, neutral). Behavioural data are presented in Figure 2-2. Significant main effects of cue were observed on RTs ($F(2,28) = 11.473$, $p < .001$, $\eta_p^2 = 0.2907$), but not on accuracy ($F(2,28) = 1.365$, $p = .272$, $\eta_p^2 = 0.047$). Both feature and spatial cues led to faster responses compared to neutral cues: spatial cue – neutral cue: $t(14) = 4.339$, $p < .001$, $d = 1.16$; feature – neutral cue: $t(14) = 4.4518$, $p < .001$, $d = 1.19$. RTs following spatial and feature cues were not significantly different from each other: $t(14) = 0.331$, $p = .745$, $d = 0.088$.

MEG session

Repeated-measures ANOVAs were carried out on accuracy and RTs from correct trials with factors of spatial cue (valid, neutral) and feature cue (valid, neutral). Behavioural data are presented in Figure 2-2. Significant main effects of spatial cue were observed on RTs ($F(1,19) = 29.562$, $p < .001$, $\eta_p^2 = 0.609$), and accuracy ($F(1,19) = 24.548$, $p < .001$, $\eta_p^2 = 0.564$). Significant main effects of feature cue were observed on RTs ($F(1,19) =$

22.902, $p < .001$, $\eta_p^2 = 0.547$) and accuracy ($F(1,19) = 4.913$, $p < .05$, $\eta_p^2 = 0.206$). No significant interactions were observed (all p's > .152).

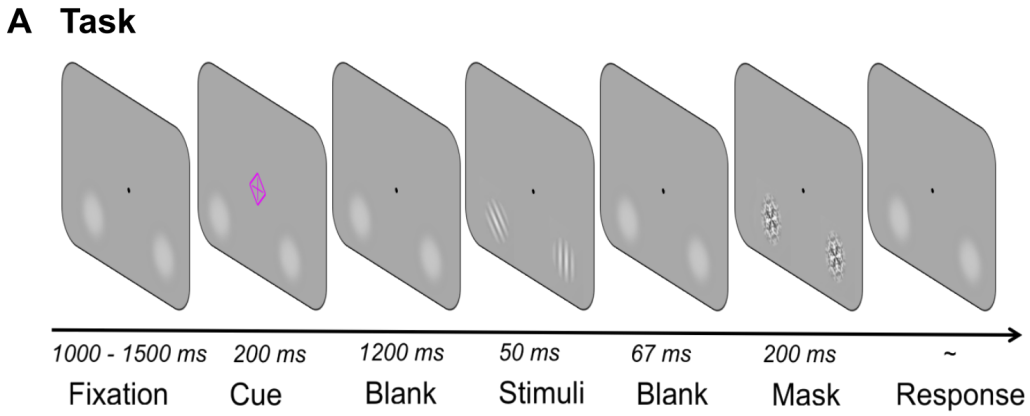

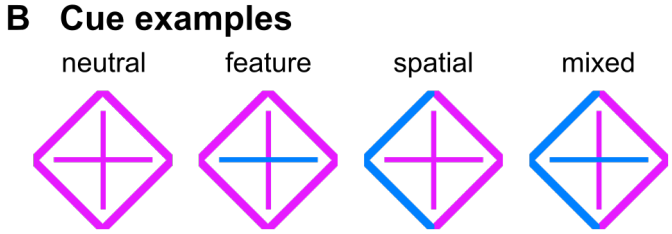

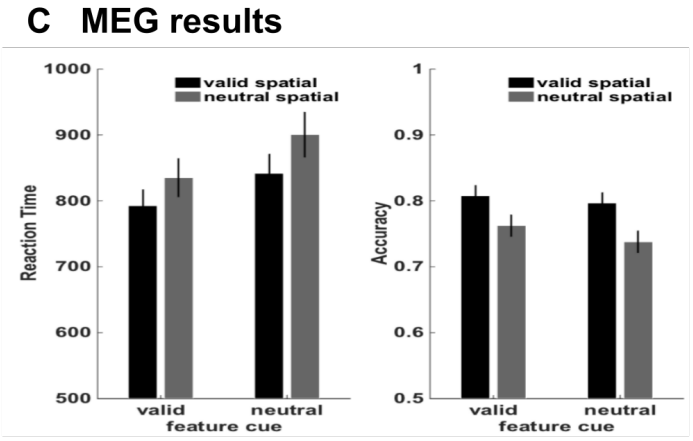

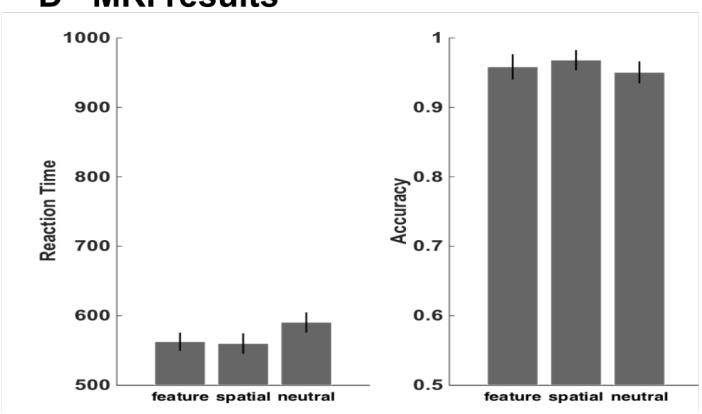

Figure 2-2. A. Trial procedure. **B.** Examples of the cues used in the different cueing conditions. Colour assignments were counterbalanced across participants. **C.** Behavioural results in the MEG session. Valid spatial cueing led to shorter RTs and higher accuracy. Valid feature cueing shortened RTs and improved accuracy. Feature and spatial cueing did not interact. **D.** Behavioural results in the MRI session. Spatial and feature cueing led to shorter RTs, but did not influence accuracy. Error bars reflect ± 1 standard error.

2.6.2 fMRI Results

I first determined which brain regions were engaged by the spatial and feature-based orienting tasks compared to rest periods (see Figure 2-3). Trial-related activity produced a marked increase of activity in occipital cortex (MOG), bilaterally, extending into the cerebellum. Frontal activations included supplementary motor area (SMA) bilaterally; regions along the precentral sulcus (PrCS) bilaterally, involving the superior frontal gryus (SFG), middle frontal gyrus (MFG) and inferior frontal gyrus (IFG); and the left anterior insula. PrCS has been proposed to be the human homolog of the frontal eye fields (FEF) (Paus, 1996: Nobre et al. 1997). The areas of parietal activation included the inferior parietal lobule, the intraparietal sulcus (IPS), and the superior parietal lobule (SPL) in the left hemisphere. Subcortical regions of activation included the basal ganglia in the putamen bilaterally and in the right caudate. These regions correspond to regions previously found in attentional orienting tasks (e.g. Shulman et al., 2002; Yantis et al., 2002).

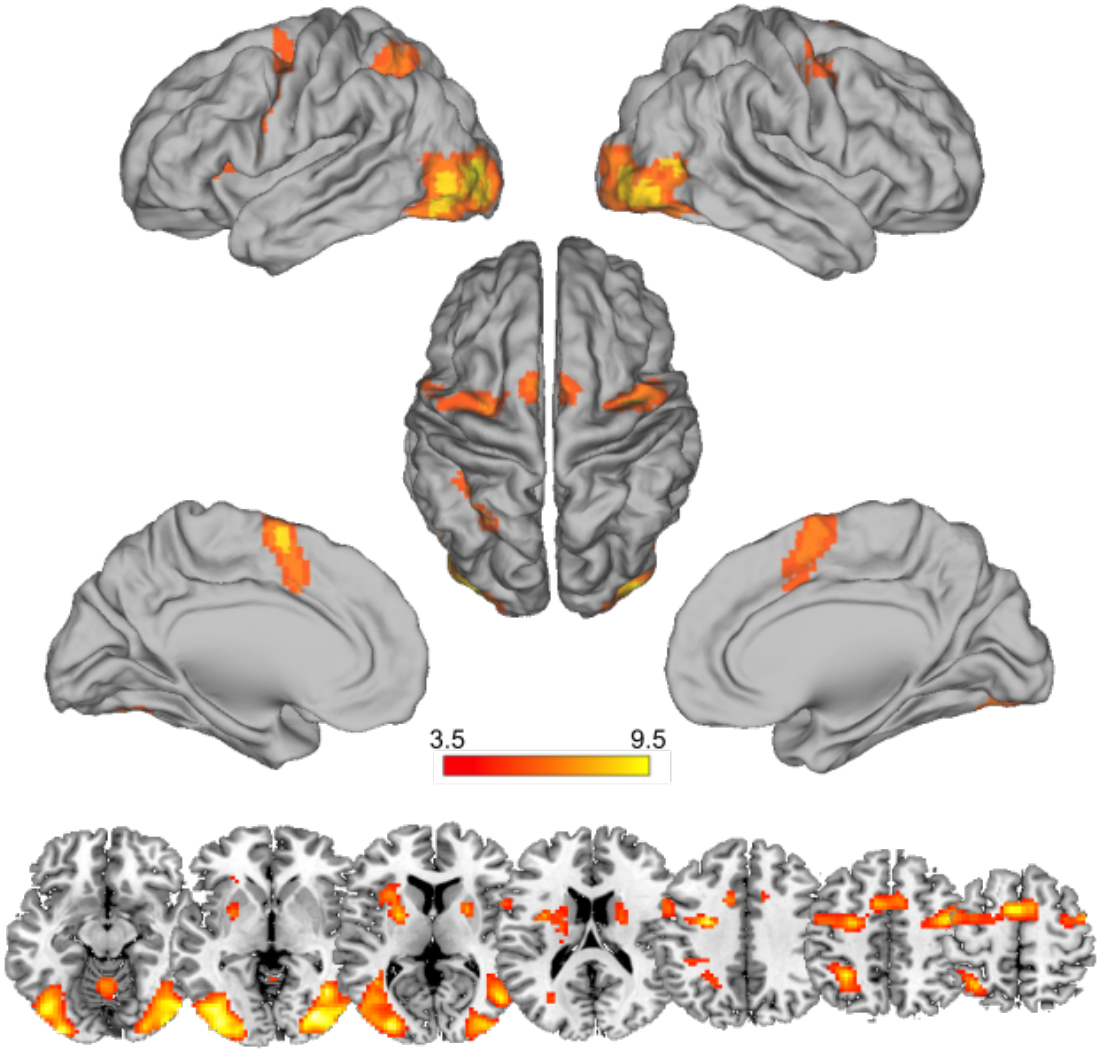

Figure 2-3. T-statistic maps for group-level task versus rest contrast. Clusters showing greater activation during the task than during rest (corrected p<.05). Colour scale indicates t values.

Trials with spatial cues compared to feature cues activated the middle temporal (MTG) gyrus, bilaterally, superior temporal gyrus (STG) in the right hemisphere, and the inferior parietal lobule in the left hemisphere. The precuneus was also preferentially activated in the left hemisphere (see Figure 2-4).

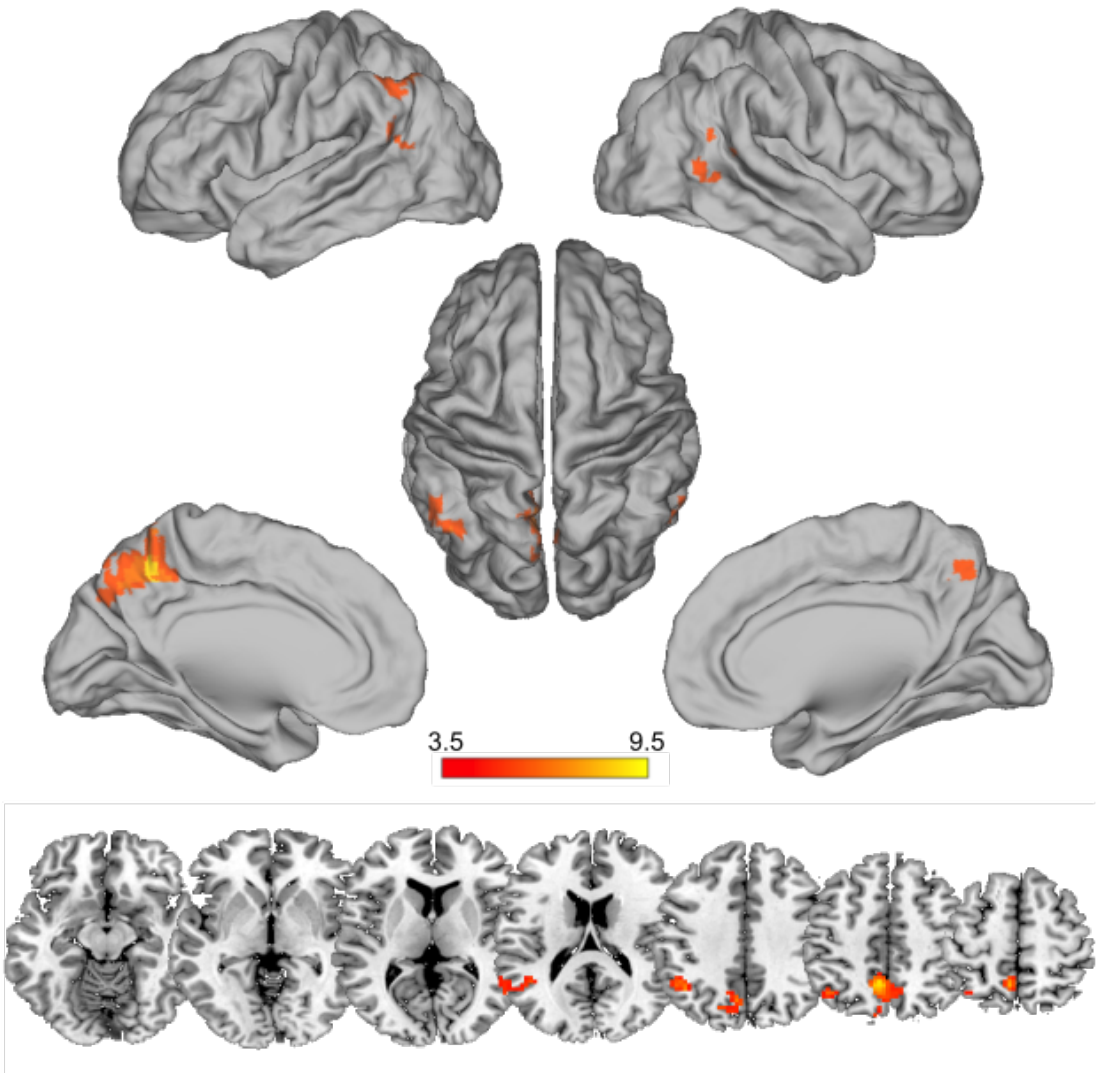

Figure 2-4. T-statistic maps for group-level spatial attention versus feature attention contrast. Clusters showing greater activation after spatial cues than after feature cues (corrected p<.05). Colour scale indicates t values.

Contrasts comparing spatial preparatory attention to general cue-related activity, contrasts comparing feature preparatory attention to general cue-related activity, and contrasts comparing feature preparatory attention to spatial preparatory attention did not reveal any significant activations.

TASK > REST				
Cortical Region of Activation	X	Y	Z	Peak Z
Left cerebellum, MOG bilaterally	33	-67	-23	5.41
Left PrCS, left insula, left putamen, left SFG, left MFG, left IFG, SMA bilaterally	-27	-7	37	4.69
Left IPL, left SPL, left precuneus	-30	-49	49	4.67
Right PrCS, right SFG, right MFG, right IFG	42	-4	49	4.60
Right putamen, right caudate	-24	2	7	4.55
SPATIAL > FEATURE ATTENTION				
Left precuneus	-6	-52	46	4.74
Right STG, right MTG	63	-49	22	4.00
Left MTG, left IPL, left angular	-45	-58	40	3.89
The Z scores are all significant at $P < 0.05$ after correcting for the whole brain.				

Table 1. Regional cortical activations for the group analysis.

2.6.3 MEG Results

Time frequency analysis: alpha-band

Effect of spatial cueing on alpha power: left vs right spatial cues

Performing a cluster permutation test on the averaged alpha-band power (8 - 14 Hz) from 0 ms to 1,400 ms revealed a significant positive cluster over left posterior sensors from 320 ms to 1,280 ms after cue onset ($p = .0044$) and a just-significant negative cluster over right posterior sensors from 440 ms to 920 ms after cue onset ($p = .0452$) (Figure 2-5). Additionally, comparing alpha power lateralisation following spatial cues without feature information versus spatial cues with feature information

(without averaging over frequency, i.e. from 3 – 30 Hz) revealed a significant positive cluster (4 Hz to 16 Hz from 440 ms to 620 ms, p = .025) (Figure 5C).

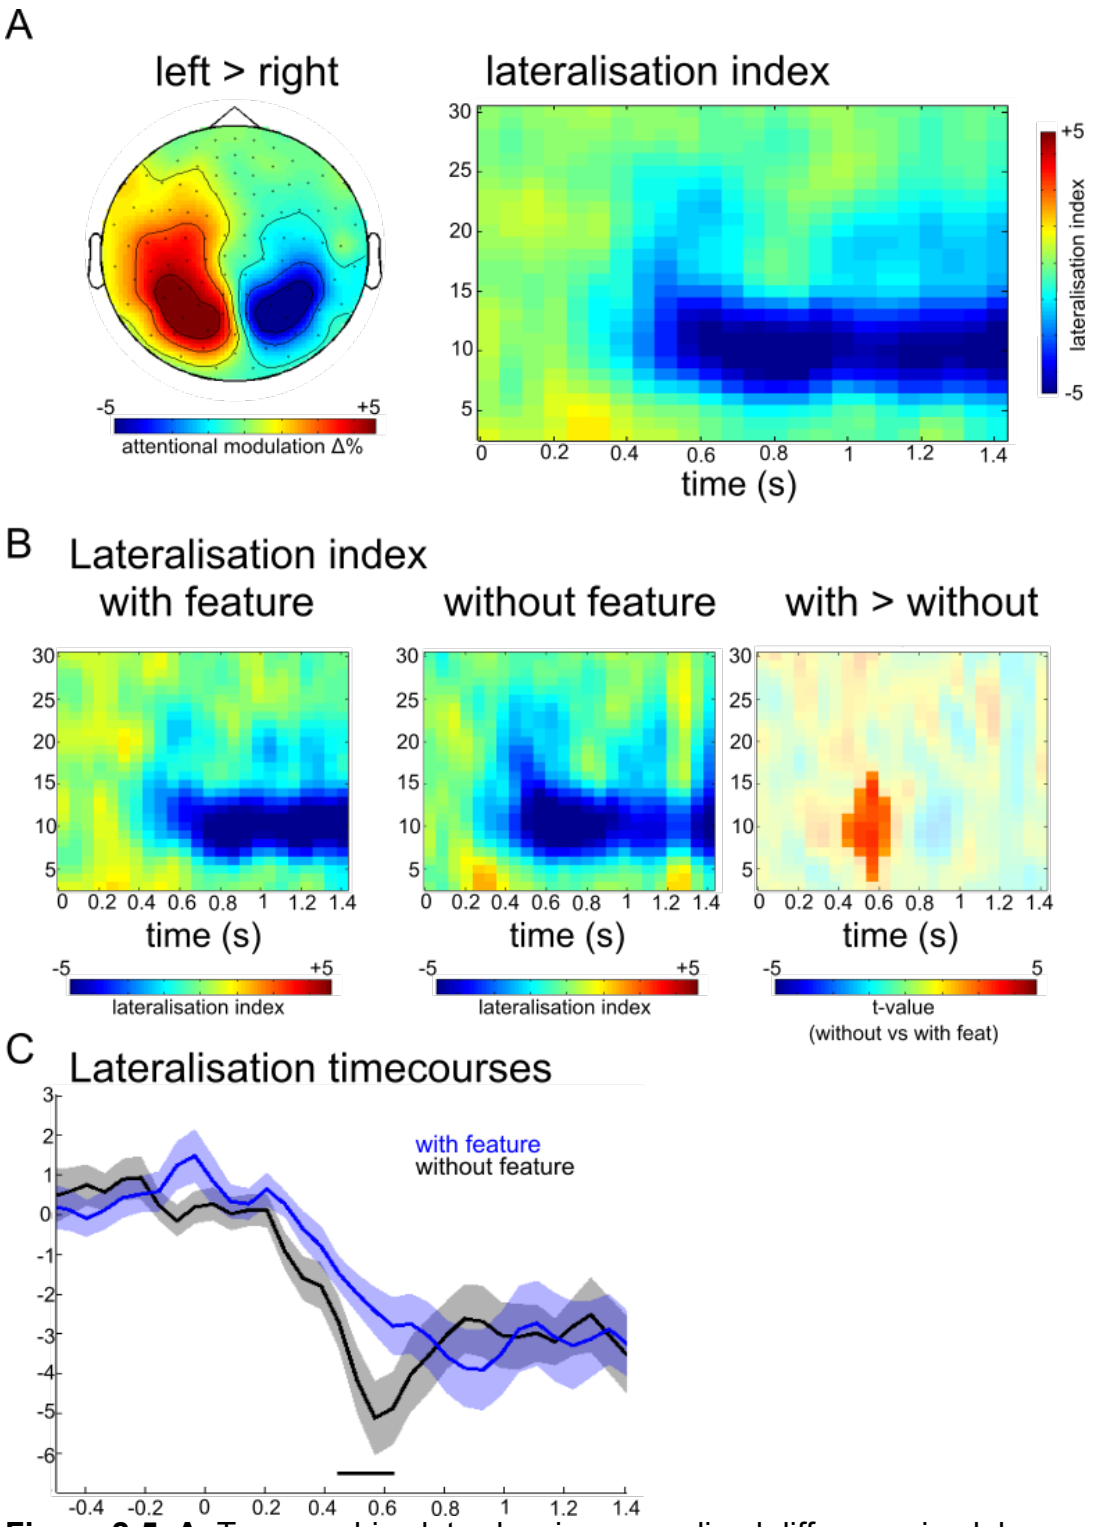

Figure 2-5. A. Topographic plots showing normalized difference in alpha power (8 – 14 Hz, averaged over significant time window) between left and

right spatial cues, expressed as percentage change (i.e., ([left - right] / [left + right]) * 100). *Right:* Average time-frequency representations of alpha lateralisation index (i.e. [(Contralateral responses - ipsilateral responses) / (contralateral responses + ipsilateral responses)]*100). Calculated for all spatial cueing types. **B**. Average time-frequency representations of alpha lateralisation indices for spatial cues with feature information, spatial cues without feature information, and the difference between the two conditions. **C**. Time courses of alpha power lateralisation index for spatial cues with and without feature information. The black bar denotes significant differences between the conditions ($p < .05$).

Effect of feature cueing on non-lateralised alpha power: feature vs non-feature cues

Performing a cluster permutation test on the averaged alpha-band power (8 - 14 Hz) from 0 ms to 1,400 ms revealed one significant negative cluster over left posterior and central sensors from 440 ms to 920 ms after cue onset ($p = .024$) (Figure 2-6A).

Feature modulation of non-lateralised spatial cueing effect: spatial cues with versus without feature information

Performing a cluster permutation test on the averaged alpha-band power (8 - 14 Hz) for left spatial cues from 0 ms to 1,400 ms revealed a significant negative cluster primarily over left but also right posterior, central and frontal sensors from 380 ms to 800 ms ($p = .001$), while a cluster permutation test for only right spatial cues revealed one significant negative cluster over right posterior, central and frontal sensors from 440 ms ms to 680 ms ($p = .011$) (Figure 2-6B and C).

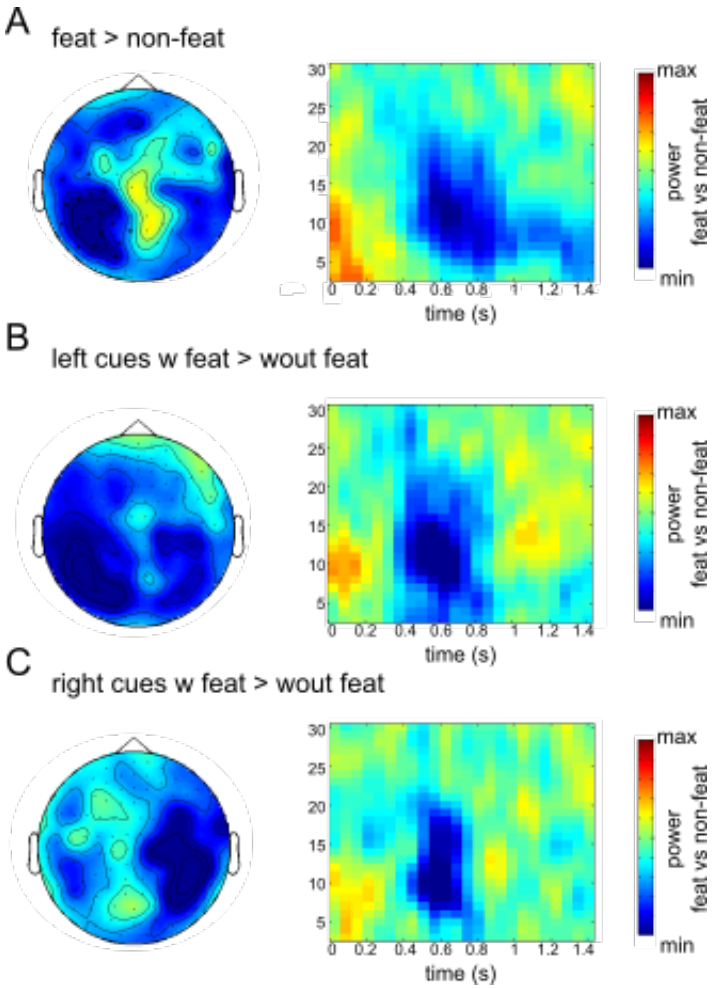

Figure 2-6. A. Topographic plots showing alpha power (8 – 14 Hz) for feature versus non-feature cues, averaged over significant time window. *Right*. Average time-frequency representations of the same contrast averaged over significant channels. **B.** Topographic plots showing alpha power (8 – 14 Hz) for left cues with versus without feature information, averaged over significant time window. *Right*. Average time-frequency representations of the same contrast averaged over significant channels. **C.** Topographic plots showing alpha power (8 – 14 Hz) for right cues with versus without feature information, averaged over significant time window. *Right*. Average time-frequency representations of the same contrast averaged over significant channels.

Time frequency analysis: gamma-band

Target-induced gamma lateralisation: spatial vs non-spatial cues

Performing a cluster-based permutation test on the averaged gamma-band power (40 - 100 Hz) from 0 ms to 600 ms after target onset revealed no significant clusters (Figure 2-7A).

Target-induced gamma lateralisation: feature vs non-feature cues

Performing a cluster permutation test on the averaged gamma-band power (40 - 100 Hz) from 0 ms to 600 ms after target onset revealed a significant positive cluster from 28 ms to 388 ms (p = .002). Running the same permutation analysis without averaging over frequency revealed a significant positive cluster at 68 Hz to 82 Hz from 28 ms to 388 ms (p = .004) (Figure 2-7B).

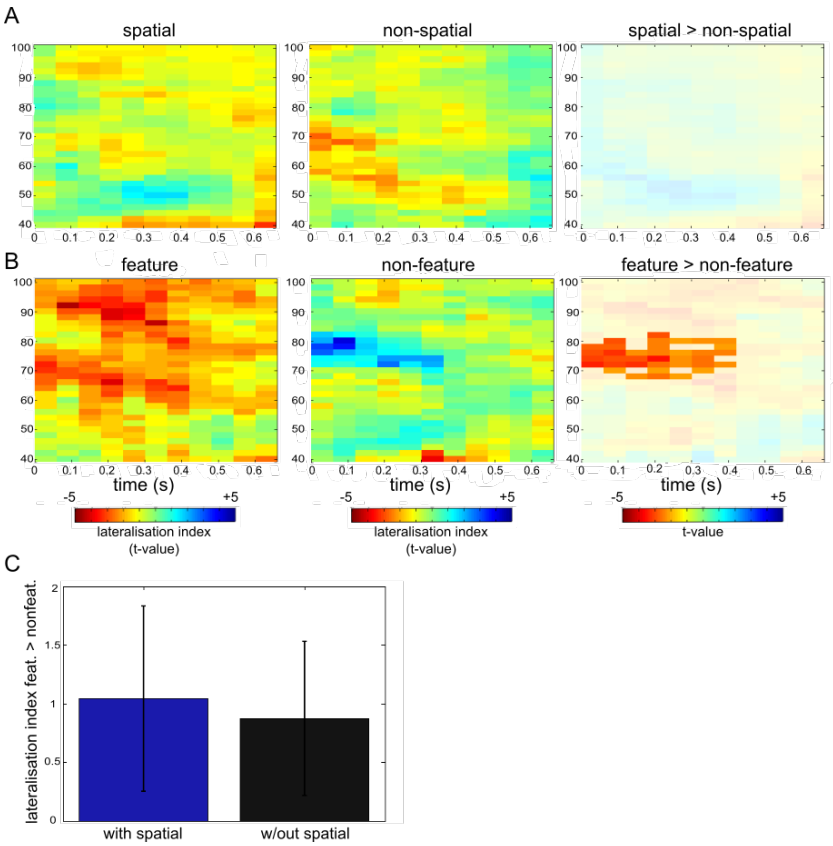

Figure 2-7. Average time-frequency representations of gamma lateralisation indices (i.e. [(Contralateral responses - ipsilateral responses) / (contralateral responses + ipsilateral responses)]*100) over central-posterior channels following **A**. spatial cues, non-spatial cues, and the difference between the conditions and **B**. feature cues, non-feature cues, and the difference between the conditions. **C**. Bar graphs showing gamma lateralisation indices for the feature versus non-feature cueing contrast, calculated separately for cues with and without spatial information (i.e., feature-spatial vs spatial cues and feature vs neutral cues), averaged over significant time window and frequencies. There was no difference between the two conditions (p > .05).

Event-related fields

Event-related fields locked to the target onset were computed for latitudinal and longitudinal gradiometers separately, and responses compared between contralateral and ipsilateral responses evoked by the targets preceded by spatial cues (both with and without feature information). In addition, contralateral and ipsilateral ERFs were separately compared for the four different cueing conditions: neutral cues, feature cues, spatial cues and combined spatial-feature cues. For latitudinal gradiometers, permutation analysis revealed significant differences between contralateral and ipsilateral responses following spatial cues ($p = .014$, from 316 ms to 348 ms post-target onset), while for longitudinal gradiometers there were no significant differences between contralateral and ipsilateral responses (Figure 2-8A and B, left). A repeated-measures ANOVA on the event-related responses in latitudinal gradiometers averaged over the significant time window (316 ms to 348 ms) with factors of feature cue (valid, neutral), spatial cue (valid, neutral) and laterality (contralateral, ipsilateral) revealed a significant spatial cue by laterality interaction ($F(1,19) = 9.680$, $p = .006$, $\eta_p^2 = 0.338$). The main effects of feature cue and laterality were also significant ($F(1,19) = 11.442$, $p = .003$, $\eta_p^2 = 0.376$ and $F(1,19) = 6.208$, $p = .022$, $\eta_p^2 = 0.246$, respectively), and there was a trend for a significant main effect of spatial cue ($F(1,19) = 3.486$, $p = .07$, $\eta_p^2 = 0.155$). (Figure 2-8A, right). The same ANOVA on longitudinal gradiometers yielded no significant results (all F's < 2) (Figure 2-8B, right).

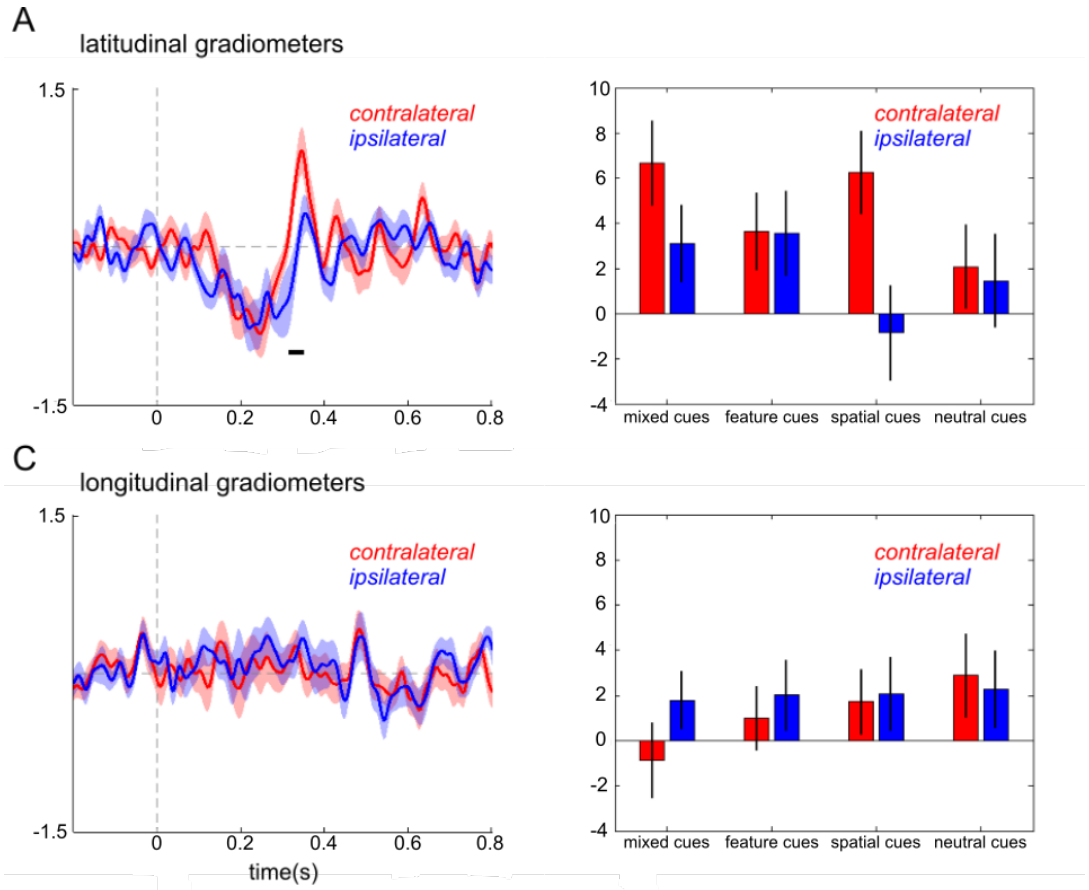

Figure 2-8. A. *Left* Contralateral and ipsilateral event-related responses to the target averaged over visual sensors following spatial cues for latitudinal gradiometers. The black bar below the ERF responses denotes significant differences between contralateral and ipsilateral responses ($p < .05$). *Right* event-related responses averaged over visual sensors following the four different cueing types, averaged over the significant time window highlighted by the black bar on the left. **B**. Same as in A but for longitudinal gradiometers. Shading indicates standard error of the mean. Y-axis indicates field strength.

2.7 Discussion

This study contrasted feature-based and spatial attention with the aim of examining how alpha and gamma power support preparation for upcoming targets and target selection on the basis of feature versus spatial information. My results suggest that feature-based and spatial attention operate largely independent but rely on similar mechanisms, consistent with previous findings (Cohen & Maunsell, 2011; Giesbrecht et al., 2003). Behaviourally, I show a benefit in performance with selective attention: feature-based and spatial attention independently facilitated behaviour. There was no evidence for an interaction between feature-based and spatial attention. MEG recordings revealed that both feature-based and spatial attention rely on alpha power modulations to bias the processing of an anticipated target, but modulations occur at different spatial scales. An enhancement of lateralised gamma responses to the target emerged when observers held feature expectations relative to when they did not, but no modulation of target-induced gamma power was observed in any other conditions.

I replicated the well-established effect of spatial cueing on alpha power: alpha power contralateral to the direction of the spatial cue was reduced relative to ipsilateral alpha power during the cue target interval. This is consistent with a large body of data (e.g. Sauseng et al., 2005; Kelly et al., 2006). This effect was only subtly altered by the presence of feature information and alpha power lateralisation was relatively delayed compared to spatial cues without feature information. However, alpha lateralisation was comparable for spatial cues with and without feature information before

and at target presentation. This delay may reflect differences in cue processing. Alternatively, it could reflect a reduction in the difference in alpha power between contralateral and ipsilateral sensors when spatial cues contain feature information, as feature information reduces alpha power bilaterally. An analysis looking at alpha-lateralisation would be blind to such an effect. I found evidence for this explanation when I examined the effects of feature information on non-lateralised alpha. First, feature cues (both with and without spatial information) lead to bilateral alpha power suppression compared to alpha power following non-feature cues (spatial cues without feature information and neutral cues), although the effect was more pronounced over left channels. Second, when I contrasted spatial cues with feature information to spatial cues without feature information I observed an additional decrease of alpha power that was most pronounced ipsilateral to the direction of spatial cues, i.e. in regions where alpha power was relatively high in pure spatial cueing trials. In contrast, I only observed modest differences between pure and mixed spatial cues on contralateral responses. This suggests the effect of feature information on spatial cues is not of an interactive nature whereby the addition of feature information enhances the spatial attention effect to further suppress alpha power contralaterally. This might be due to a ceiling effect were alpha power already maximally suppressed in contralateral regions. Instead, the addition of feature information interacts with spatial cues by primarily modulating ipsilateral responses. In other words, mixed feature-spatial cues lead to a relatively more global alpha power suppression extending over both hemispheres compared to pure spatial cues, which modulated alpha power

locally. A modulation of alpha power by feature information was also observed independently from spatial cues: alpha power was significantly reduced for feature cues compared to non-feature cues over left posterior regions.

Together, these findings support the notion that spatial attention modulates alpha power locally in a retinotopic fashion, while feature attention modulates alpha power globally across the whole visual hemifield. They extend our understanding of top-down modulations based on feature information: previous demonstrations of alpha power modulations by feature-based attention contrasted attentional shifts from one feature dimension to another (e.g. from colour to motion, Snyder et al., 2010) or contrasted trials where a colour patch was attended versus trials where it was not (Müller and Keil, 2004). However, it is unclear how top-down attention mechanisms can target different features from within the same feature dimension, which are represented in overlapping brain areas and brain cells. These results demonstrate that alpha-band power modulations are relatively flexible and can operate on the basis of spatial and feature information. This mechanism does not depend on retinotopic or topographic representations and can also target groups of spatially-dispersed neurons.

I did not find an increase in gamma power to stimulus presentation with spatial attention (e.g. Wyart et al., 2008; Koelewijn et al., 2013; Landau et al., 2007, Bauer et al., 2006). This discrepancy might be for a number of reasons. First of all, target presentation on its own did not elicit an increase in gamma power, which most likely reflects the stimulus choice in my design. The gamma response is very sensitive to stimulus properties and is

strongest at maximum contrast levels, a spatial frequency of 3 cpd and for stimuli that extend across several degrees of the visual field (Adjamian et al., 2004; Hall et al., 2005; Henrie & Shapley, 2005). For example, (Hall et al., 2005) showed that the amplitude of gamma oscillations measured using MEG increased linearly with stimulus contrast. The stimuli used in my study were staircased to threshold and the average stimulus contrast across participants was around 15% (±3%) Michelson contrast, a level at which only a relatively weak gamma response was observed in the Hall study. The stimuli used in this study may not have been ideal for inducing visual gamma responses as they were small and presented at low contrast. Indeed, other studies that successfully showed stimulus-induced gamma using MEG used large and high-contrast stimuli (Bauer et al., 2014; Gruber et al., 1999; Koelewijn et al., 2013; Landau et al., 2007; Siegel et al., 2008). Wyart et al. (2008) reported stimulus-induced gamma responses with small, low-contrast stimuli but this study presented stimuli for a much longer duration (400 ms) compared to 50 ms in the present study. In short, while I did not find an effect of spatial attention on target-induced gamma, I do not claim that these effects do not exist but rather that the experimental paradigm employed here was not sensitive enough.

I did, however, observe a modulation of target-induced gamma lateralisation by whether or not observers held feature expectations during target presentation: lateralised gamma power was significantly increased in trials with feature cues compared to trials without feature cues. This finding is broadly consistent with the Match and Utilisation Model (MUM, Herrmann et al. 2004). According to this model early (i.e. before 150 ms after stimulus

onset) gamma responses reflect the match between memory content and bottom-up input with larger responses when there is a match than when there is not. Early gamma responses are dependent on both top-down and bottom-up information and gamma activity is involved in signalling matches. Importantly, later gamma responses (after 200 ms after stimulus onset) to the target are assigned a different function. Later gamma responses are proposed to 'utilise' the match information to initiate and guide downstream processes.

My results are consistent with the MUM: in the feature cueing condition participants likely have to maintain feature information in memory and, in target present trials, this information matches bottom-up input. According to the MUM, I should observe a relatively larger gamma response in this condition compared to conditions where top-down expectations do not match bottom–up input, or if they match to a lesser degree. In the neutral cueing condition observers either have no feature expectations, or maintain two feature expectations simultaneously, depending on the task strategy they adopt. In both those scenarios the match between bottom-up and top-down information is not as complete as in feature cueing trials. The timing of the feature cueing effect on gamma power is also consistent with the MUM: gamma power modulation was already significant at 28 ms after stimulus onset. This effect persisted until 388 ms after target onset. According to the MUM the effect could, therefore, reflect two separate processes: the matching of bottom-up input and top-down expectations. I only presented sensor-space data but whether there are distinct sources of the gamma power at different latencies could be tested in source-space.

Furthermore, the MUM account also predicts that a gamma power response should only emerge in target-present trials but not in target-absent trials. It was not possible to assess the effect of target presentation on the feature modulation of gamma power as gamma lateralisation was calculated relative to the side of target presentation[5].

Gamma activity has also been proposed to signal prediction error, i.e. when bottom-up input is unexpected and does not match expectations (Arnal & Giraud, 2012). This account, however, predicts increased gamma power when there is a mismatch between bottom-up and top-down input. This proposal draws support from findings that gamma power scales with prediction error, mismatch negativity and expectation violations (Arnal, Wyart, & Giraud, 2011; Axmacher et al., 2010; Gruber & Müller, 2005; Haenschel, Baldeweg, Croft, Whittington, & Gruzelier, 2000; Nicol et al., 2012; Todorovic, van Ede, Maris, & de Lange, 2011). However, this account does not make a distinction between early and late gamma effects. Under the MUM only early gamma power, before 150 ms post-stimulus onset, reflects the degree of matching between memory contents and bottom-up input while later gamma effects are linked to utilisation of the matching information. Importantly, all the reported increases in gamma power following expectation violation occur outside the 'early' time window proposed by Herrmann et al. (2004). The majority of expectation effects on

[5] For the feature vs nonfeature contrast, gamma lateralisation had to be calculated relative to the target side as the nonfeature cueing condition includes neutral cueing trials. Notably, when I compared gamma lateralisation relative to cue direction (only including target absent trials) for spatial-feature cues vs spatial cues there was no difference in gamma responses to the probe display, in line with the predictions made by the MUM.

gamma were reported between ~200 ms to ~ 400 ms post stimulus (Arnal et al., 2011; Gruber & Müller, 2005; Haenschel et al., 2000; Todorovic et al., 2011) some as late as 1000 ms post stimulus (Axmacher et al., 2010). The earliest effect reported was by Nicol et al. (2012) who reported that unexpected stimulus changes increased long-range gamma synchronization between frontal, temporal, parietal and occipital cortical areas 128 to 256 ms after stimulus onset. Therefore, these results do not necessarily conflict with my results and may instead reflect later mechanisms also supported by gamma power.

The results from the fMRI session were mostly consistent with previous findings. The task-related activity pattern is consistent with the frontoparietal control network (e.g. Shulmann et al., 2002; Yantis et al. 2002). Only few regions were more activated following spatial cues compared to feature cues including the MTG, bilaterally, the STG in the right hemisphere, the IPL, the angular gyrus and the precuneus, all in the left hemisphere. Giesbrecht et al. (2003) reported selective activation following spatial cues in left SFG / MFG, in left SPL and the right IPL. They also reported selective activation for feature cues in fusiform gyrus, bilaterally, extending into the inferior temporal gyrus (ITG). Schenkluhn et al. (2008) used transcranial magnetic stimulation (TMS) to assess the causal contribution of different parietal regions in feature-based and spatial attention. Stimulation of the anterior intraparietal sulcus impaired preparatory spatial- and feature-based selection, whereas TMS of the supramarginal gyrus (SMG) impaired spatial selection only. Stimulation of the posterior intraparietal sulcus did not disrupt either process.

These results are substantially different to ours. Firstly, I did not observe any feature-specific activation. Secondly, the only region that was activated selectively following spatial cues in both my and Giesbrecht's study was the IPL but in different hemifields. However, given the large differences in task design, and the fact my task design was not optimal for uncovering subtle differences as I used a blocked design with relatively long blocks, it is not surprising to see some differences between my results and previous results. The general pattern of activity is consistent with previous reports and show that both feature-based and spatial attention rely on a similar, highly-overlapping network of brain areas.

Using behavioural, MEG and fMRI methodology I show that feature-based and spatial attention are independent: I did not observe interactions at the behavioural or the neural level. My results suggest that both rely on similar mechanisms: a common frontoparietal network was activated for spatial and feature cues that only showed subtle differences between cueing conditions. Anticipatory biasing was supported by alpha power modulations for both feature-based and spatial attention but this was implemented in a different manner for selective attention on the basis of feature information versus spatial information. Feature-based attention modulated brain activity in the alpha-band globally across both hemispheres while spatial attention operated locally. In addition, my findings also suggest a role of gamma power in matching memory contents against bottom-up input: when observers carried feature expectations target-induced gamma lateralisation was significantly greater relative to when observers did not carry detailed feature expectations. Together, the findings extend our

understanding of the functioning of feature-based and spatial attention as functionally distinct mechanisms in the flexible biasing of information processing.

3 Subliminal and Supraliminal Bottom-up Influences on WM Versus Perceptual Representations

3.1 Chapter Abstract

Effects of supraliminal distracters on WM representations and WM performance are well established, but the underlying changes that lead to the impairments observed in performance are not well characterized. Similarly, while we know that low-signal information of which observers are not consciously aware can nevertheless influence perceptual processes, it is not clear whether such subliminal information might exert an influence on WM representations. Relatively few studies have examined the interactions between subliminal and supraliminal information in WM. In this chapter, I present three experiments examining these two questions: can subliminal stimuli influence WM performance? How are WM representations affected by the presentation of distracting information? First, in Experiments 1a and 1b I replicate the finding that orientation stimuli can influence behaviour subliminally in a visuomotor priming task. Experiments 2 and 3 used the same orientation stimuli, but participants had to remember a target orientation and report it back by adjusting a probe orientation after a memory delay. Before or after presentation of the target orientation, a subliminal or supraliminal distracter orientation was presented that was either irrelevant for task completion and never had to be reported (Experiment 2), or was relevant for task completion because it had to be reported on some trials (Experiment 3). In both experiments, presentation of a supraliminal distracter influenced WM recall of the target orientation.

When the distracter was presented subliminally, however, there was no impact on orientation recall. These results suggest that information stored in WM is protected from influences of subliminal stimuli, while online information processing is modulated by subliminal information.

3.2 Introduction

Subliminal stimuli have been shown to influence processing of subsequent stimuli at a number of levels - from low-level visual priming up to and including semantic priming (Dehaene et al., 1998; Eimer & Schlaghecken, 1998; 2002; Greenwald, Draine, & Abrams, 1996; Neumann & Klotz, 1994). Acceptance of subliminal effects is not universal (Pratte & Rouder, 2009; Reingold, 2004), but there are numerous demonstrations of subliminal influences on response times in priming tasks (Kiefer et al., 2011; Kouider & Dehaene, 2007). For example, Eimer and Schlaghecken (2002) examined the influence of masked prime stimuli - which could not be identified above chance - on responses to a target stimulus and showed that both subliminal and supraliminal prime stimuli influenced participants' RTs to target stimuli though interestingly in different ways. These results were replicated and their interpretation supported by additional control studies (see Eimer and Schlaghecken, 2003, for a review).

Past studies have focused on the effects of subliminal stimuli on responses to stimuli immediately present in the observer's environment. However it remains unknown whether subliminal information can influence what becomes consciously available in WM. WM representations are short-lived representations that are strictly limited in their capacity (Curtis & D'Esposito, 2003; Zhang & Luck, 2008). Theorists have conceptualised WM

in different ways. For example, one common conception is that WM depends on attentional selection of stimuli into a limited capacity store (Baddeley, 1986; Bundesen, 1990), which is distinct from perceptual and semantic representations of the stimuli. Alternatively, it has been argued that WM reflects the temporary activation of perceptual and semantic representations of stimuli (Cowan, 1995).

These different theories posit different putative effects of subliminal processing on WM. For example, if WM involves the activation of perceptual and semantic representations of stimuli, then there is little reason to think that WM should be immune from influence of subliminal information. The same perceptual and semantic representations would be recruited in WM identification tasks that use similar stimuli to tasks in which visuomotor priming has been established. On the other hand, if WM is abstracted from the perceptual and semantic representations used on-line for object identification, then it may be that stimuli must be supraliminal, entering the limited-capacity WM store, to influence processing.

Silvanto and Soto (2012) examined the influence of a subliminal visual distracter on performance in a WM task. Participants had to retain a target orientation in memory to compare against a probe orientation after a delay period. During the delay period a masked distracter orientation was presented on some trials. Accuracy in the task was significantly reduced when an incongruent distracter was presented during the WM delay period, relative to both congruent and no-distracter conditions, suggesting that distracting subliminal stimuli influence the maintenance and/or retrieval of WM representations.

One important aspect of the Silvanto and Soto study is that they used subjective rather than objective measures of subliminal processing, which may have underestimated the degree of conscious access to the masked stimulus (see Hannula, Simons, and Cohen, 2005, for a review). In a control experiment, forced-choice identification was very high (d' > .80), suggesting that the stimuli may have been available to objective report, although they were not available to subjective, conscious report. This difficulty can be overcome if objective measures of awareness are taken as well as subjective measures. Whereas subjective measures are good for evaluating an observer's internal, subjective experience of the environment, they might be based not only on sensitivity to stimulus presence but also on the observer's decision criterion. Observers may indicate no awareness when their perceptual experience of a stimulus simply failed to surpass their criterion for confidently reporting it.

Here I present three experiments testing effects of subliminal and supraliminal stimuli on perceptual and WM reports. In all cases objective and subjective measures of stimulus awareness were used. The first experiment was an adaptation of studies by Eimer and Schlaghecken (1998; 2002) using stimulus parameters matched to the subsequent WM experiments. In Eimer and Schlaghecken (1998), prime and target stimuli were assigned either the same, or opposite, responses (on congruent and incongruent trials). The prime stimulus, when masked, influenced responses to the target: RTs were slower and more errors were made following congruent compared to both neutral and incongruent primes. In addition, neural markers of motor preparation (the lateralized readiness potential,

LRP, measured using event-related potentials; ERPs) were also affected by prime-target congruency. Specifically, prime stimuli first elicited their corresponding response but this initial activation then reversed, meaning an incongruent response was activated for congruent trials while the congruent response was activated for incongruent trials. The authors argued that the effects reflect inhibition of the response initially activated by the prime, when the prime is masked (Eimer and Schlaghecken, 1998; 2002). I aimed to replicate the behavioural results using the stimuli employed in the other experiments here, which targeted WM rather than perceptual report. Two versions of the experiment were conducted, using blocked (Experiment 1a) and trial-by-trial (Experiment 1b) measures of subjective and objective awareness.

Experiment 2 was a WM task in which I asked participants to remember the orientation of one of two sequentially presented gratings for later recall. One orientation was a target and the other a distracter. I varied the visibility of the stimuli, rendering one subliminal in some conditions (based on objective measures) and supraliminal in others. I examined whether subliminal information could influence WM representations. Experiment 3 was a replication and extension of Experiment 2. I assessed whether any influence of subliminal information might depend on its task relevance by occasionally probing the subliminal stimuli so that they became task-relevant.

3.3 General Methods

3.3.1 Participants

All experimental protocols were reviewed and approved by the Central University Research Ethics Committee of the University of Oxford. A total of 85 volunteers participated in the experiments: 31 males and 54 females, ranging from 18 to 34 years of age. All participants had normal or corrected-to-normal vision and were naïve to the purpose of the experiment. No participants participated in more than one experiment. Participants gave informed consent before taking part, and received course credits or financial compensation (£8 per hour) for taking part.

3.3.2 Experimental Procedure

Experiments were prepared and presented using the Psychophysics Toolbox in Matlab (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007). The stimuli were presented on a 23-inch light-emitting-diode (LED) display with a spatial resolution of 1920x1080 and a vertical refresh rate of 120 Hz. Stimuli were presented in dark grey (14.5 cd/m2) at the centre of the screen against a light grey background (25.7 cd/m2). Orientation stimuli used throughout the study consisted of the contours of an oriented bar (2.5° visual angle in length and 0.5° visual angle in width) superimposed on a circle (1° visual angle in diameter) forming the outline of a 'UFO' shape (Figure 3-1a). The probe orientation stimulus was identical to the distracter and target orientation stimulus except for its colour (light green, 21.3 cd/m2). This change in colour made the probe item easily distinguishable from the to-be-remembered orientation stimuli in the sequence. Pattern masks were

created on each trial by overlaying the contours of 20 black (1.7 cd/m2) randomly oriented bars, each with a randomly determined offset of maximally 1° visual angle from the centre of the screen (max. length and width: 3.5° visual angle). Thus, the exact appearance of the mask varied from trial to trial. The bars used in creating the mask were identical to those used to create the 'UFO' stimuli.

3.4 Experiment 1a: Influence of subliminal stimuli in a visual priming task

In Experiment 1, I aimed to demonstrate that distracting orientation stimuli, which remain subliminal by objective criteria, influence orientation judgements about an orientation probe stimulus. To this end, I set out to replicate priming effects previously reported by Eimer and Schlaghecken (1998, 2002), but here using orientation, rather than arrow, stimuli. Eimer and Schlaghecken (1998, 2002) reported opposite effects on RTs following subliminal and supraliminal primes. Specifically, congruent supraliminal primes speeded RTs to the probe. When primes were rendered subliminal this effect reversed, and RTs were faster following incongruent compared to congruent primes. I predicted the same pattern of effects for my experiment.

3.4.1 Methods

Participants

Twenty volunteers took part (7 male, 18 – 30 years, one left-handed).

Stimuli

Following the design of Eimer & Schaghecken (1998, 2002), two orientations were used throughout the task: left (120°) and right (60°). Two different types of masks were employed: a dense and a sparse mask. The

dense mask was created as described under general methods. The sparse mask consisted of only 5 randomly overlapping horizontal, vertical, and oblique lines (length = 2.5° visual angle, width = 0.12° visual angle) each with a randomly determined offset of maximally 0.36° visual angle from the centre (see Figure 3-1a for an example of stimuli and masks used). The dense mask was intended to render stimuli subliminal, while the sparse mask was intended to preserve stimulus visibility (Eimer and Schlaghecken, 1998, 2002).

Task and Procedure

The experiment comprised two tasks: A visual priming task and a prime-identification task. The purpose of the priming task was to examine whether the prime orientation stimuli could influence RTs to the subsequent probe orientation stimuli, in other words whether these kinds of subliminal stimuli are sufficient to influence behaviour. The purpose of the identification task was to ensure that prime stimuli were rendered subliminal by the dense mask but remained supraliminal with the sparse mask. Figure 3-1a provides a graphical summary of the priming task. Each trial began with the presentation of a central fixation cross (500-ms duration). A prime stimulus (16-ms duration) was immediately followed by a mask (70-ms duration), then followed by a 50-ms blank screen, and finally by the target stimulus (100-ms duration). Participants responded according to the orientation of the target (left or right). Visual feedback (‘correct, ‘incorrect’, ‘too slow’) was shown for 500 ms after a response was made or after 450 ms (whichever occurred first). The inter-trial interval was 1,300 ms. Prime and target orientations were either the same, or opposite (i.e., congruent or

incongruent). In other words, if the prime stimulus was 60° the target stimulus could be 60° or 120°. Both congruency conditions were equiprobable and randomized within each block. Mask type was varied between blocks, and block order was counterbalanced across participants. Participants were instructed to respond as quickly as possible without sacrificing accuracy, using their left and right index fingers for ‘left’ and ‘right’ orientations, respectively.

The priming task consisted of a factorial design with two within-subject factors: probe (same, different) and visibility (supraliminal, subliminal). Participants completed eight blocks of 40 trials each (320 trials in total, 80 trials per condition). Four blocks were run using the dense mask (160 trials), the other four with the sparse mask (160 trials). Blocks alternated between dense-mask blocks and sparse-mask blocks, and the starting block-type was counterbalanced across participants. Breaks were given after every block. Participants completed a practice block of 40 trials before they started the priming task.

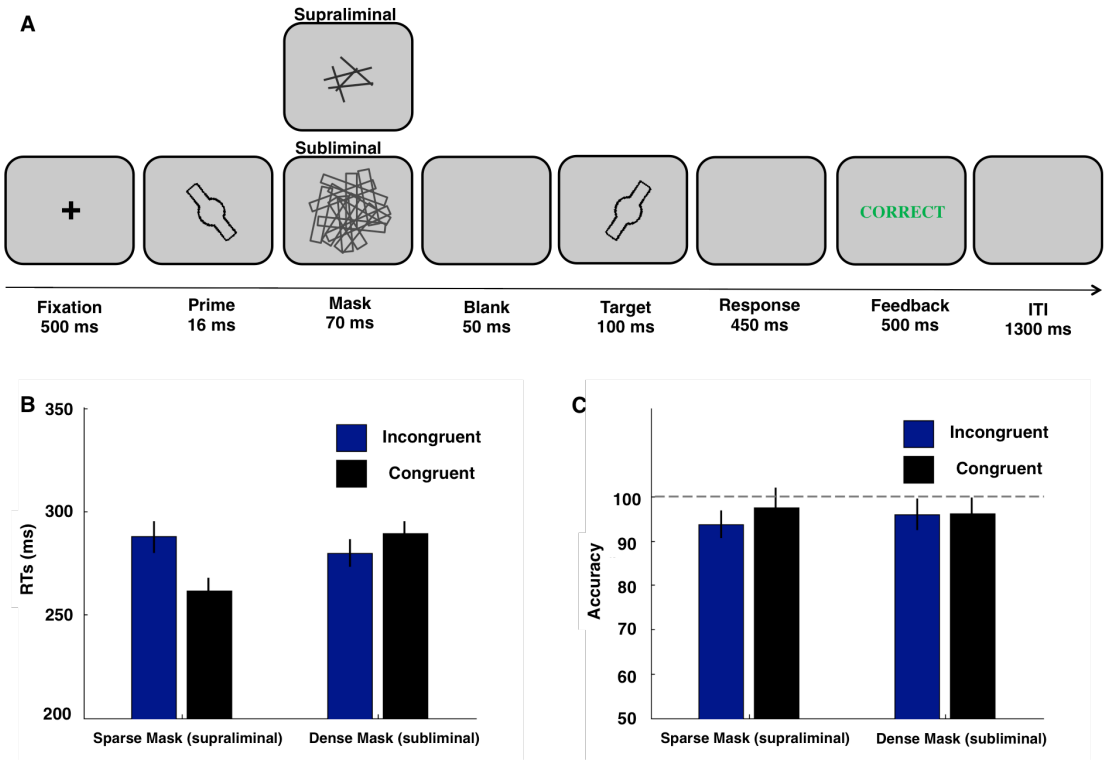

Figure 3-1. Design for 1a and 1b, and results for 1a of the visual priming task. **A**, A central fixation cross was presented for 500 ms followed by a 16-ms prime stimulus, and then a 70-ms dense or sparse backwards mask. After a 50-ms mask-target interval, a target stimulus (congruent or incongruent with the prime stimulus) appeared for 100 ms. Participants had 450 ms to indicate the orientation (left or right) of the target stimulus and were given 500 ms feedback on their performance. After a 1,300-ms intertrial interval the next trial began. **B**, Median reaction times to probes following congruent and incongruent prime stimuli followed by a dense or sparse mask (subliminal and supraliminal, respectively) in Experiment 1a. Error bars reflect ± 1 standard error. **C**, Mean accuracy to probes following congruent and incongruent prime stimuli followed by a dense or sparse mask (subliminal and supraliminal, respectively) in experiment 1a. Error bars reflect ± 1 standard error.

A visual identification task was used to obtain objective and subjective measures of awareness for stimuli followed by dense and sparse masks (see Figure 3-2a). The same stimuli and masks as in the priming task were used. Participants were presented with a stimulus in half the trials to assess observers' ability to discriminate between stimulus presence and absence as well as their subjective awareness of the stimulus. When

present, the stimulus appeared for 16 ms and was followed by a 70-ms presentation of either the dense or sparse mask. Mask density was varied between blocks, and block order was counterbalanced across participants. Following a delay of 500 ms, a probe stimulus was shown (congruent or incongruent with equal probability). In the other half of the trials, when no stimulus appeared, the procedure was the same except that a blank display instead of a stimulus was shown before the mask. Participants first completed a forced-choice discrimination task indicating whether the probe was of the same or different orientation as the stimulus. Next, participants rated their subjective awareness using the Perceptual Awareness Scale (PAS, Ramsøy, Rams, & Overgaard, 2004). This scale consists of four response options defined as follows: 1 – stimulus not seen, 2 – weak glimpse: something was there but I do not know its orientation, 3 – almost clear image: I think I know the orientation, and 4 – clear image. Participants were instructed on how to use this scale before the first session of the identification task. It was emphasized that ratings are to be made introspectively, relying on visual experience (Ramsøy & Overgaard, 2004). A shortened version of the four ratings was spelled out on the screen at the end of each trial before ratings were made (1 – stimulus not seen, 2 – weak glimpse, 3 – almost clear image, 4 – clear image). Participants responded by pressing the corresponding number on the keyboard.

The identification task consisted of the same 2x2 (probe: same, different; visibility: supraliminal, subliminal) design as the priming task with the additional factor of prime presence (present, absent). Thus, when no prime was present, probe orientation cannot be described as “same” or

"different" relative to the probe, but instead was randomly determined to be either 120° or 60°. Participants completed four blocks of 50 trials each (200 trials in total). Thus, there were a total of 25 trials per target-present condition. Two blocks were run with the dense mask (100 trials). The other two were run with the sparse mask. Blocks alternated between dense-mask and sparse-mask blocks, and the starting block-type was counterbalanced across participants.

Participants completed one experimental session lasting approximately 45 minutes. All participants first completed the priming task and then completed the identification task.

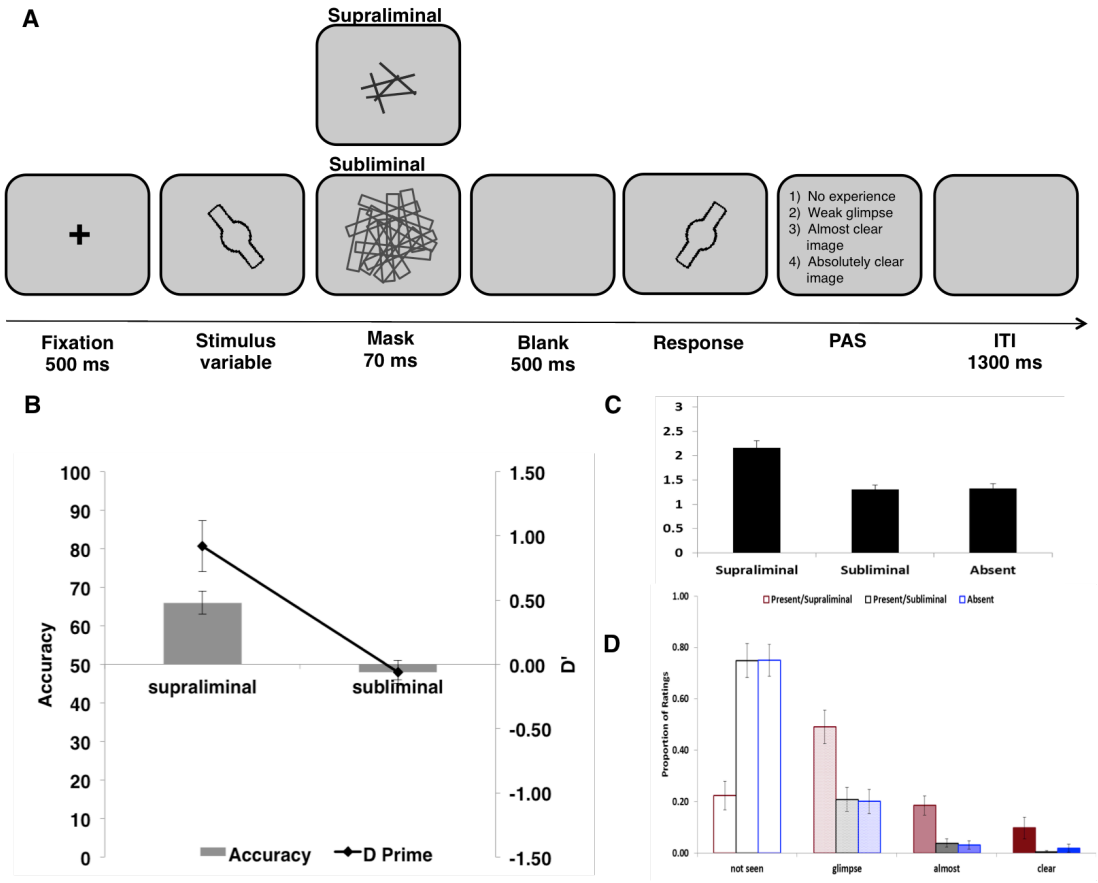

Figure 3-2. Design of the visual identification task used in Experiment 1a and 1b, and results of the visual identification task Experiment 1a. **A**, A

central fixation cross was presented for 500 ms followed by a 16 ms prime stimulus (present on 50% of trials, blank shown on other 50% of trials), and then a 70 ms dense or sparse backwards mask. After a 50 ms mask – probe interval a probe stimulus (congruent or incongruent with the prime stimulus) appeared for 100 ms. Participants first completed a forced-choice discrimination task indicating whether the probe was of the same or different orientation as the stimulus. Next, participants rated their subjective awareness using the Perceptual Awareness Scale (PAS, Ramsøy & Overgaard, 2004). After a 1,300 ms intertrial interval the next trial began. **B**, Accuracy and d' in the forced-choiced identification task for the dense and sparse mask condition separately. **C**, Mean awareness ratings for the sparse-mask, dense-mask, and stimulus absent condition separately. **D**, Proportion of different awareness ratings in the identification task for the three conditions. Error bars reflect ± 1 standard error.

3.4.2 Analysis

Priming Task

For the priming task, I analysed median RTs and accuracy. For RTs, I calculated the median rather than the mean since it is more robust to outliers. Only correct responses were included in the RT analysis.

Identification Task

Performance in the identification task was measured in levels of accuracy, d' in stimulus-present trials, and subjective ratings of awareness for both stimulus-present and stimulus-absent trials. D' was calculated from the hit rate and false-alarm rate using the following equation (Macmillan & Creelman, 1991):

$$d' = z(Hit) - z(FA)$$

Hit rate was defined as the probability that the participant responded "different" when the probe orientation was different, and the false-alarm rate was defined as the probability that the participant responded "different" when the probe orientation was the same.

In the case of hit rates of 1 and/or false alarm rates of 0, values were adjusted using the following equation (Macmillan & Creelman, 1991, p.8):

Adjusted Hit =1-1/(2*Nr of Trials)

Adjusted FA =1/(2*Nr of Trials)

3.4.3 Results

Visual Priming Task

Figure 3-1b and 3-1c show RTs and accuracy for each condition. I replicated the expected pattern of speeded RTs following congruent supraliminal primes and slowed RTs following congruent subliminal primes (Eimer and Schlaghecken, 1998, 2002). This was reflected in a highly significant mask-by-congruency interaction ($F(1,19) = 45.32$, $p < .001$, $\eta^2 = .71$). To examine this interaction further, Bonferroni-corrected paired t-tests were conducted comparing congruent and incongruent trials for each mask condition separately. The difference between responses in congruent and incongruent trials was significant for both the sparse and dense masks. In the sparse-mask condition, responses were significantly faster following congruent primes compared to incongruent primes ($M_{Congruent}$ = 261 ms ± 6, $M_{Incongruent}$ = 288 ms ± 8; $t(19) = 4.76$, $p < .001$, $d = 1.06$). In the dense-mask condition, this effect reversed, and responses were significantly faster following incongruent primes compared to congruent primes ($M_{Congruent}$ = 289 ms ± 6, $M_{Incongruent}$ = 280 ms ± 7; $t(19) = -2.43$, $p = .025$, $d = 0.54$). The main effects of probe congruency and mask were also significant ($F(1,19) = 4.48$, $p < .05$, $\eta^2 = .19$, and $F(1,19) = 5.04$, $p < .05$, $\eta^2 = .21$, respectively). The same pattern of results was observed with mean, rather than median, RTs (Mask-by-congruency interaction: $F(1,19) = 51.23$, $p < .001$, $\eta^2 = .73$;

main effect of congruency: $F(1,19) = 5.12$, $p < .05$, $\eta^2 = .21$; main effect of visibility: $F(1,19) = 4.96$, $p < .05$, $\eta^2 = .207$; t-tests dense-mask condition: $M_{Congruent} = 290$ ms ± 6, $M_{Incongruent} = 282$ ms ± 7; $t(19) = -2.36$, $p = .029$, $d = 0.53$; sparse-mask condition: $M_{Congruent} = 264$ ms ± 6, $M_{Incongruent} = 288$ms ± 8; $t(19) = 4.98$, $p < .0001$, $d = 1.11$).

For accuracy, there was a significant mask-by-congruency interaction ($F(1,19) = 10.26$, $p < .01$, $\eta^2 = .35$). Bonferroni-corrected paired-samples t-tests, comparing congruent and incongruent trials for each mask condition separately, revealed that accuracy was significantly better following congruent primes compared to incongruent primes in the sparse mask condition ($M_{Congruent} = 98\% \pm .6$, $M_{Incongruent} = 94\% \pm .8$, $t(19) = -4.08$, $p < .01$, $d = 0.91$). In the dense-mask condition, there was no difference in accuracy between congruent and incongruent prime trials ($M_{Congruent} = 96\% \pm .7$, $M_{Incongruent} = 96\% \pm .7$, $t(19) = -.193$, ns). The main effect of congruency was also significant ($F(1,19) = 9.29$, $p < .01$, $\eta^2 = .33$).

Prime-Identification Task with Blocked Design

Data from one participant in the prime-identification task were not available as the data files were mistakenly overwritten.

Accuracy and d'

Figure 3-2b shows accuracy and d' in the forced-choice discrimination task for the stimulus present/supraliminal and present/subliminal condition separately. A paired-samples t-test compared accuracy in stimulus-present conditions using dense versus sparse masks. There was a significant difference between the conditions reflecting that participants performed significantly better in the sparse mask condition than

in the dense-mask condition ($t(18) = -4.78$, $p < .001$, $d = 1.50$, M_{sparse} = 66% ± 3, M_{dense} = 48% ±2). Furthermore, performance in the dense-mask condition was not significantly better than chance ($t(18) = -.722$, ns).

Similarly, d' in the sparse-mask and dense-mask conditions differed significantly ($t(18) = 4.64$, $p < .001$, $d = 1.48$). Participants performed significantly better in the sparse-mask than in the dense-mask condition (M_{sparse} = 0.92 ±0.2, M_{dense} = -.06 ±.09). Importantly, d' in the dense-mask condition was not significantly different from 0 ($t(18) = .476$, ns). These results show that participants were not better than chance at detecting and discriminating orientation in the dense-mask condition. However, stimuli followed by the sparse mask were reliably detected and their orientation discriminated.

Subjective Ratings

Figures 3-2c and 3-2d show the mean awareness ratings and proportion of different ratings for each condition separately. A Wilcoxon signed-rank test comparing average awareness ratings in the sparse-mask, dense-mask, and absent conditions indicated that there was no difference in ratings between the absent (M = 1.32 ± 0.10) and dense-mask condition (M = 1.30 ± 0.09, Z = -0.26, ns, see Figure 3-2). However, there was a significant difference between the average awareness ratings in the sparse-mask (M = 2.16 ± 0.15) and absent condition (Z = -3.82, $p < .001$), and the sparse-mask and dense-mask conditions (Z = -3.82, $p < .001$).

To examine what was driving these differences, I calculated the proportions of each rating for each participant separately, and then ran eight separate Wilcoxon Signed-ranks tests comparing the proportions of ratings

between the absent- and sparse-mask condition, and the dense-mask and sparse-mask condition for each rating option separately. “Not Seen” ratings occurred significantly more often in the dense-mask and absent condition compared to the sparse-mask condition ($Z = -3.77$, $p < .001$ and $Z = -3.82$, $p < .001$, respectively). Conversely, “Weak Glimpse”, “Almost Clear” and “Absolutely Clear” ratings were made significantly more often in the sparse-mask condition compared to the absent and dense-mask condition (all p’s < 0.01).

Together these results suggest that participants were unable to discriminate stimulus orientations on dense-mask trials, and unable to differentiate reliably between dense-mask and stimulus absent trials. Stimulus orientation on sparse-mask trials, however, was reliably discriminated.

3.4.4 Interim Discussion

The results replicated previous findings (Eimer and Schlaghecken, 1998, 2002), showing that both subliminal and supraliminal orientation stimuli can influence discrimination responses to a probe stimulus. Participants were faster and more accurate to respond to an orientation stimulus when a supraliminal prime of the same orientation preceded the target stimulus. When the prime stimulus was presented subliminally, this effect reversed and participants were faster to respond to an orientation stimulus when the prime orientation was different to the target orientation. In contrast to supraliminal primes, there was no effect of subliminal primes on accuracy. These results show that both orientation stimuli (as used in this

experiment), and arrow stimuli (as used in Eimer and Schlaghecken, 1998, 2002), can influence behaviour when presented subliminally.

3.5 Experiment 1b: Replication of visual priming with alternative assessment of prime awareness

In Experiment 1a, I used a blocked design to test for objective awareness of stimuli in the prime-identification task. However, it has been shown that blocked designs can underestimate prime awareness (Pratte & Rouder, 2009), as participants may become disengaged and unmotivated in difficult, subliminal blocks. When subliminal and supraliminal trials are mixed together, relatively easy and relatively difficult trials vary randomly keeping participants engaged with the task. Correspondingly, estimates for chance identification of subliminal primes are different in mixed compared to blocked identification tasks. To ensure that the effects of subliminal primes in Experiment 1a were not due to residual awareness of the stimuli, I tested another 20 participants on the same visuomotor priming task followed by a prime-identification task in which sparse and dense mask conditions were randomly intermixed on a trial-by-trial basis.

3.5.1 Methods

Participants

Twenty volunteers took part (9 male, 18 – 34 years old, one left-handed).

Stimuli, Task, Procedure, & Analysis

The design, parameters, and analysis of the visual priming task were identical to Experiment 1a. Experiment 1b was therefore a straight replication of the previous experiment on an independent set of participants.

The prime-identification task was also equivalent, with the exception that the dense-mask and sparse-mask trials were intermixed randomly, on a trial-by-trial basis, within a single testing block.

3.5.2 Results

Visual Priming Task

I replicated the pattern of results observed in Experiment 1a. A significant mask-by-congruency interaction in median RTs ($F(1,19) = 42.06$, $p < .001$, $\eta^2 = .69$) showed significant and opposite effects of prime congruency in the sparse-mask ($M_{Congruent}$ = 246 ms ± 7, $M_{Incongruent}$ = 266 ms ± 7; $t(19) = 4.64$, $p < .001$, $d = 1.04$) and dense-mask conditions ($M_{Congruent}$ = 276 ms ± 7, $M_{Incongruent}$ = 270 ms ± 8; $t(19) = -2.02$, $p = .058$, $d = 0.45$). As before, main effects of congruency and visibility were also significant ($F(1,19) = 4.44$, $p = .048$, $\eta^2 = .181$, and $F(1,19) = 14.94$, $p = .001$, $\eta^2 = .44$, respectively). The same pattern of results was observed with mean RTs (Mask-by-congruency interaction: $F(1,19) = 43.53$, $p < .001$, $\eta^2 = .56$; main effect of congruency: $F(1,19) = 4.41$, $p = .049$, $\eta^2 = .188$; main effect of visibility: $F(1,19) = 13.57$, $p < .01$, $\eta^2 = .42$; sparse-mask condition: $M_{Congruent}$ = 251 ms ± 7, $M_{Incongruent}$ = 268 ms ± 6; $t(19) = 4.39$, $p < .001$, $d = 0.98$; dense-mask condition: $M_{Congruent}$ = 277 ms ± 7, $M_{Incongruent}$ = 272 ms ± 6; $t(19) = -2.63$, $p = .016$, $d = 0.58$).

Similarly, I replicated the findings in accuracy. A significant mask-by-congruency interaction ($F(1,19) = 34.49$, $p < .001$, $\eta^2 = .65$) revealed significantly better accuracy following congruent primes compared to incongruent primes in the sparse-mask condition ($M_{Congruent}$ = 97% ± .7, $M_{Incongruent}$ = 92% ± 1.4; $t(19) = -4.57$, $p < .01$, $d = 1.02$) and no difference in

the dense-mask condition ($M_{Congruent}$ = 95% ± 1.6, $M_{Incongruent}$ = 95% ± 1.1; t(19) = 0.25, ns). The main effect of congruency was also significant ($F(1,19) = 7.70$, $p = .012$, $\eta^2 = .29$).

Analysis of performance in the prime-identification task revealed that five participants performed better than chance at identifying the subliminal prime (see below). Importantly, when I excluded those participants from analysis, I observed the same pattern of results. A significant mask-by-congruency interaction in median RT data ($F(1,14) = 35.47$, $p < .001$, $\eta^2 = .72$) showed opposite effects of congruency in the sparse-mask ($M_{Congruent}$ = 244 ms ± 8, $M_{Incongruent}$ = 259 ms ± 8; $t(14) = 3.17$, $p = .007$, $d = 0.81$) and the dense-mask ($M_{Congruent}$ = 273 ms ± 8, $M_{Incongruent}$ = 267 ms ± 9; $t(14) = -2.10$, $p = .054$, $d = 0.54$) conditions. A significant mask-by-congruency interaction in the accuracy data ($F(1,14) = 19.78$, $p = .001$, $\eta^2 = .59$) was driven by a significant effect of congruency in the sparse-mask condition ($M_{Congruent}$ = 97% ± .9, $M_{Incongruent}$ = 92% ± 1.7; $t(14) = -3.32$, $p < .01$, $d = 0.86$) but not in the dense-mask condition ($M_{Congruent}$ = 94% ± 1.3, $M_{Incongruent}$ = 94% ± 2.1; $t(14) = -0.19$, ns).

Prime Identification Task with Mixed Design

Accuracy and d'

A paired-samples t-test compared accuracy in stimulus-present conditions using dense versus sparse masks. There was a significant difference between the conditions reflecting that participants performed significantly better in the sparse mask condition than in the dense-mask condition ($t(19) = -4.67$, $p < .001$, $d = 1.04$, M_{sparse} = 65% ±4, M_{dense} = 49%

±2). Furthermore, performance in the dense-mask condition was not significantly better than chance (t(19) = - .39, ns).

Similarly, d' in the sparse-mask and dense-mask conditions differed significantly (t(19) = -3.54, p <.01, d = 0.79). Participants performed significantly better in the sparse-mask than in the dense-mask condition (M_{sparse} = 1.03 ±0.33, M_{dense} = - .04 ±0.11). Importantly, d' in the dense-mask condition was not significantly different from 0 (t(19) = .68, ns). These results show that, at the group level, participants were not better than chance at detecting and discriminating orientation in the dense-mask condition. However, stimuli followed by the sparse mask were reliably detected and their orientation discriminated. At the individual subject level five participants performed significantly better than chance (significant binomial test on accuracy data at the .05 level), suggesting they could differentiate between present and absent trials to some extent.

Subjective Ratings

I also replicated the pattern of subjective ratings reported in Experiment 1a. A Wilcoxon signed-rank test comparing average awareness ratings in the sparse-mask, dense-mask, and absent conditions indicated that there was no difference in ratings between the absent (M = 1.46 ± 0.14) and dense-mask condition (M = 1.42 ± 0.13, Z = -0.48, ns). However, there was a significant difference between the average awareness ratings in the sparse-mask (M = 2.19 ± 0.15) and absent condition (Z = -3.77, p < .001), and the sparse-mask and dense-mask conditions (Z = -3.77, p < .001, respectively).

To examine what was driving these differences, I calculated the proportions of each rating for each participant separately, and then ran eight separate Wilcoxon Signed-ranks tests comparing the proportions of ratings between the absent- and sparse-mask condition, and the dense-mask and sparse-mask condition for each rating option separately. “Not Seen” ratings occurred significantly more often in the dense-mask and absent condition compared to the sparse-mask condition ($Z = -3.73$, $p < .001$ and $Z = -3.92$, $p < .001$, respectively). Conversely, “Weak Glimpse”, “Almost Clear” and “Absolutely Clear” ratings were made significantly more often in the sparse-mask condition compared to the absent and dense-mask condition (all p’s < 0.04).

Together these results suggest that participants were unable to discriminate stimulus orientations on dense-mask trials, and unable to differentiate reliably between dense-mask and stimulus absent trials. Stimulus orientation on sparse-mask trials, however, was reliably discriminated. When I excluded the 5 participants who performed better than chance at the identification task from analysis of subjective ratings the same pattern of results was observed.

3.5.3 Interim Discussion

In Experiment 1b, I used a mixed design to assess prime identification. This revealed some awareness for the prime stimuli in a small number of participants in my sample. This highlights the importance of considering motivational factors when assessing observers’ awareness in a subliminal prime-identification task (Pratte & Rouder, 2009). Precautions must be taken to ensure observers stay engaged with the task, for example

by mixing subliminal with supraliminal trials. Most importantly, however, I replicated the main findings of subliminal priming in the visuomotor priming task. The replication of the visuomotor priming effects (Eimer and Schlaghecken, 1998, 2002) in two independent sets of participants provides reassurance about their reliability, even though the magnitude of the effects in each case is modest.

Under the Bayesian framework, it is possible to combine data sets and as one adds more data one is simply adding more evidence to prior (however vague) knowledge (Dienes, 2011). Capitalizing on the identical design and procedures of the two priming tasks, I pooled the data across the two experiments (without excluding any participants) and computed the JZS Bayes factor on the comparison of median (mean) RTs and accuracy between congruent and incongruent primes in the dense-mask condition. For the combined median RTs (N = 40; incongruent = 275 ms; congruent = 283 ms; t = -3.18) the JZS Bayes Factor = 11.97; and for combined mean RTs (N = 40; incongruent = 277 ms; congruent = 284 ms; t = -3.42), the JZS Bayes factor = 21.54. That is, the data strongly favour the alternative over the null hypothesis, suggesting that subliminal prime stimuli influence behaviour response times in the task used here. When the 5 participants from Experiment 1b who showed residual awareness were excluded I observed similar values (median RTs = 11.3; mean RTs = 23.65).

3.6 Experiment 2: Influence of task-irrelevant subliminal vs supraliminal stimuli on WM

Having replicated the finding that subliminal stimuli can influence visuomotor processes in a priming task, I aimed to examine whether

subliminal stimuli can influence stimuli that become consciously available in WM representations. Participants were presented with a sequence of two orientation stimuli – one of which was followed by a mask. Participants recalled the orientation of the unmasked, target, stimulus. The masked stimulus acted as a distracter. On half the trials, the distracter stimulus was subliminal and on the other half it was supraliminal. I examined whether recalling an orientation from WM can be influenced by a distracter orientation, and to test how this varies with changes in the observer's awareness of the distracter orientation. In line with the visuomotor priming design by Eimer and Schlaghecken (1998, 2002), I was mainly interested in examining the influence of a distracter before encoding, rather than during the delay period.

3.6.1 Methods

Participants

Twenty-four new volunteers took part in the experiment (5 male, 18 – 30 years old, all right-handed). Two additional participants started but did not complete all sessions of the experiment. Their data were not included in the analysis.

Stimuli

The orientations of target, distracter, and probe stimuli were randomly determined on each trial and ranged from 1 to 180°. Distracter duration and presentation was varied to yield three conditions occurring in equal proportions: distracter absent, supraliminal distracter present (presented for 200 ms), and subliminal distracter present (duration individually determined, M = 36 ms ± 7; see below for further details). When

the distracter was absent, a blank image was shown for either the duration used in the supraliminal or the subliminal condition (randomly determined on each trial). The duration of the mask was 70 ms.

Task

The experiment comprised two tasks: A WM task (see Figure 3-3A) and a visual identification task (see Figure 3-4A). The purpose of the WM task was to examine whether the orientation of the masked distracter stimulus could influence encoding of the target orientation. The purpose of the identification task was to ensure the mask rendered the distracting stimulus indiscriminable in the subliminal condition.

In the WM task, participants were instructed to report the orientation of the unmasked stimulus. Each WM trial began with the presentation of a central fixation cross for 500 ms followed by either the target (unmasked) or distracter (masked) orientation stimulus. Since I was primarily interested in testing the effect of subliminal and supraliminal stimuli on WM encoding, the proportion of trials in which the distracter stimulus appeared first was larger (80% of the trials). To ensure that participants attended to all stimuli in the sequence, the distracter was presented after the target on the remaining 20% of the trials. There was a 500-ms delay between offset of the first orientation stimulus (or the subsequent mask if the first orientation stimulus was the distracter) and onset of the second orientation stimulus. After another 500-ms delay after the second orientation stimulus or mask, participants reported the target orientation by rotating a probe stimulus to the remembered target orientation by moving the computer mouse.

The prime identification task was identical to the task used in Experiment 1, except that only the dense-mask, but not the sparse-mask, condition was included (see Figure 3-4A). The probe stimulus was of the same orientation as the stimulus, or an orientation changed by 90° to the left or right (equally likely).

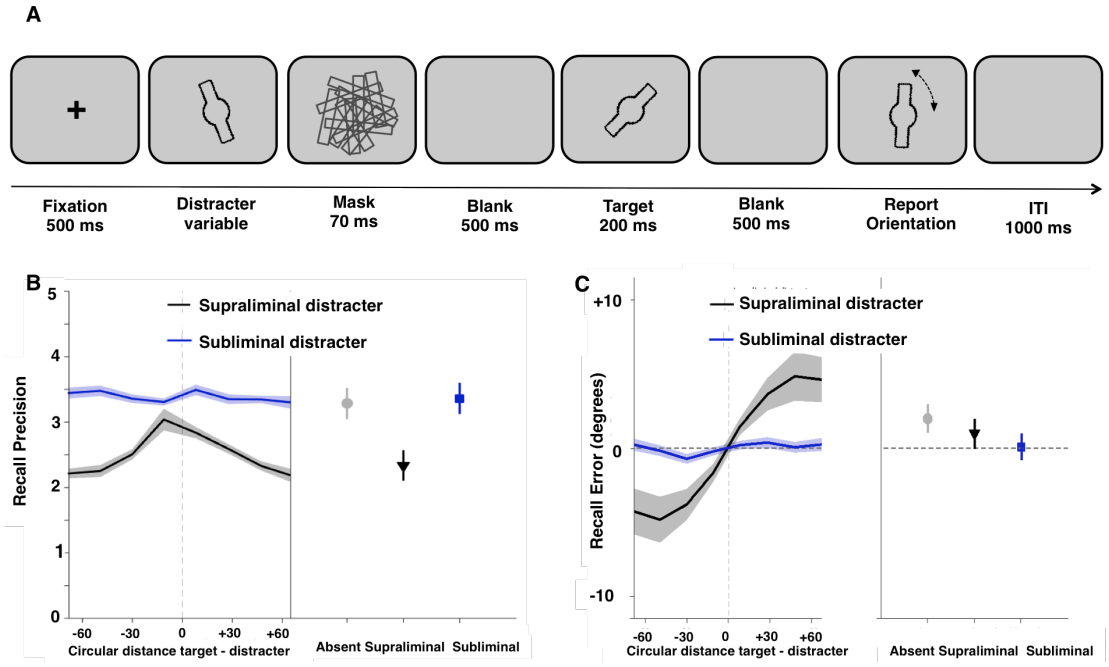

Figure 3-3. Design and results of the WM task used in Experiment 2. **A**, A central fixation cross was presented for 500 ms followed by either the target or distracter orientation stimulus (distracter first on 80% of trials). Participants were instructed to report the orientation of the unmasked stimulus. Thus, the target was always unmasked and the distracter was always immediately followed by a mask. There was a 500-ms delay between offset of the first orientation stimulus (or the subsequent mask if the first orientation stimulus was the distracter) and onset of the second orientation stimulus. After another 500-ms delay, participants reported the target orientation by rotating a probe stimulus to the remembered target orientation by moving the computer mouse. **B**, Mean recall precision (1/SD) for the distracter conditions as a function of target – distracter similarity (left) and mean overall recall precision for the two conditions (N = 23) in Experiment 2. For plotting purposes, the data for this and all similar figures depicting response error and recall precision were smoothed by using 8 overlapping bins in the analysis where adjacent bins overlap by 50% and each bin contains 25% of the overall data. Shaded areas and error bars reflect ±1 standard error. **C**, Mean recall error for the supraliminal and

subliminal conditions as a function of target-distracter similarity (left) and mean overall recall error for absent, supraliminal, and subliminal conditions (N = 23) in Experiment 2.

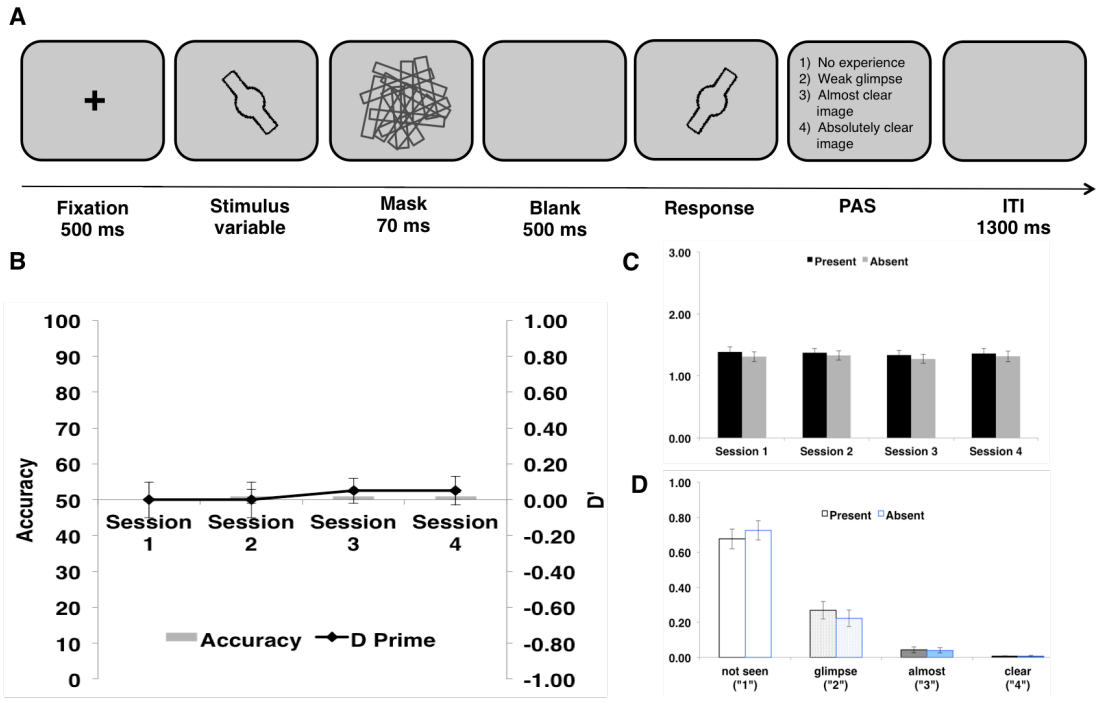

Figure 3-4. Design of the prime identification task used in Experiments 2 and 3, and results of prime identification task in Experiment 2. **A**, The task sequence in the prime identification task was identical to Experiment 1 with the exception that the dense-mask, but not the sparse-mask condition, was condition was included. **B**, Accuracy and d' in the forced-choiced identification task for stimulus present condition as a function of session number in Experiment 2. Error bars reflect ±1 standard error of the mean. **C**, Mean awareness ratings for the stimulus present and stimulus absent condition in Experiment 2 plotted for each session separately. **D**, Proportion of awareness ratings in the identification task for stimulus-present and stimulus-absent condition.

Design and Procedure

Participants completed two sessions on two separate days; the first session lasted approximately two hours and the second session lasted approximately 90 minutes. Before the beginning of the first session, participants completed a staircase procedure to set the duration of stimulus presentation for the subliminal condition. After the staircase procedure,

participants completed a practice session of 50 trials in which they received feedback after every trial. The participants then started the WM task. In each session, participants completed both the WM and the identification task, always starting with the WM task. The WM task consisted of a factorial design with the variables “distracter” (absent, supraliminal, subliminal) and “stimulus order” (target first, distracter first). The different trial types were randomly intermixed within a block. Only trials in which the distracter appeared first were used in the analysis.

Overall, participants completed a total of 1200 trials (24 blocks of 50 trials) of the WM task, resulting in a total of 320 trials in each condition in which the distracter appeared before the target, and 80 trials in each condition in which the distracter appeared after the target. The first 600 trials were completed in the first session, the second 600 trials in the second session. Participants completed a total of 400 trials (four blocks of 100 trials) of the identification task, resulting in a total of 25 trials per condition in each block. The prime-identification task was completed after every sixth block of the WM task (every 300 trials).

Staircase Procedure

The distracter duration in the subliminal condition was determined before the start of the experiment using a staircase procedure. The staircase task used the same stimulus sequence as the identification task, but without subjective ratings (see Figure 3-4A). Participants completed 60 trials of the staircase task. Stimuli and masks were identical to those in the WM task. Each trial started with the presentation of an orientation stimulus, followed by a 70-ms mask, and then by a delay of 500 ms, and finally the

presentation of a probe stimulus. Participants indicated whether the probe stimulus was identical or rotated by 90° relative to the first orientation. Probe stimuli were of the same orientation in half the trials. I used a one-up/one-down staircase to derive the presentation duration of the first orientation that lead to 50% accuracy. The staircase started with a presentation duration of 50 ms, and dynamically adapted the duration in log steps based on participants' accuracy every 20 trials. Participants received feedback at the end of each trial in the form of a written word ("correct, "incorrect") presented centrally on the screen. The staircase procedure was repeated if necessary until participants reliably performed at chance level.

3.6.2 Analysis

Working-Memory Task

Only trials in which the distracter appeared first were included in the analysis. For each trial, the reported orientation was collected and measured against the orientation of the target stimulus to derive recall error and recall precision. Recall error was defined as the angular deviation between the target orientation and the reported orientation. Recall precision measured the trial-to-trial variability in the response error, and was calculated as the reciprocal of the standard deviation (SD) of errors across trials. Since orientation space is circular, I used Fisher's definition of SD for circular data (Fisher, 1993), and subtracted the value expected by chance so that a precision value of zero corresponds to chance performance. In a one-item WM task, recall precision is estimated around 3.5 radian-1, and this value decreases as number of items is increased (Bays et al., 2009). First, I compared overall precision in the three conditions. Then, to examine

whether, and how, presentation of a distracter orientation had an effect on information stored in WM I estimated recall precision and recall error as a function of target-distracter similarity – i.e. the angular difference between the target and distracter orientations. Data from each participant were sorted into 8 equally sized non-overlapping bins (n ≈ 40 trials per bin per subject), ranging from -90° to + 90° angular deviation. Recall precision and error were then calculated separately for each subject and condition for each of the 8 bins.

Hypotheses regarding the effects of experimental parameters on recall precision and recall error were tested using ANOVAs and t tests.

Identification Task

Performance in the identification task was measured as described in Experiment 1.

3.6.3 Results

Working-Memory Task

Recall precision

Overall precision was significantly different between the subliminal-distracter, supraliminal-distracter, and distracter-absent conditions ($F(2,46) = 28.70$, $p < .001$, $\eta^2 = 1.11$, see Figure 3-3), reflecting that precision in the supraliminal-distracter condition was significantly worse than in both the subliminal-distracter and the distracter-absent conditions ($M_{Supraliminal} = 2.34$ radian-1 ± 0.221, $M_{Subliminal} = 3.36$ radian-1 ± 0.228, and $M_{Absent} = 3.29$ radian-1 ± 0.226, respectively. $t(23) = -5.75$, $p = < .001$, $d = 1.17$ and $t(23) = -5.34$, $p = < .001$, $d = 1.09$). Recall precision in the subliminal-distracter and distracter-absent conditions were not different ($t(23) = 1.124$, ns).

To compare the effects of subliminal vs. supraliminal stimuli on WM encoding, I used 2-by-8 repeated-measures ANOVAs with the variables "visibility" (supraliminal, subliminal) and "target-distracter similarity" (8 non-overlapping bins ranging from -90° to + 90° angular deviation) on the recall precision and recall error data. Figure 3-3 shows the recall precision with which participants recalled the target's orientation as a function of target-distracter similarity for each visibility condition separately. There was a significant interaction between visibility and target-distracter similarity ($F(7,161) = 4.28$, $p < .001$, $\eta^2 = .157$). To examine this interaction further, two separate one-way repeated-measures ANOVAs assessed the effect of target-distracter similarity for each condition separately. There was a significant effect of target-distracter similarity in the supraliminal ($F(7,161) = 6.00$, $p < .001$) but not in the subliminal ($F(7,161) = 0.49$, ns) condition. Specifically, the more similar the orientations of the distracter and the target were, the better participants' recall precision tended to be (see Figure 3-3B).

Recall error

Figure 3-3C shows recall error for the target orientation as a function of target-distracter similarity for each visibility condition separately. The mean error in the distracter-absent condition was 0.59° ± 0.45°, which was not significantly different from zero ($t(23) = 1.34$, ns). Visual inspection of the graph suggests that recall error varied as a function of target-distracter similarity in the supraliminal but not the subliminal-distracter condition. When the distracter orientation was clockwise to the target orientation (a negative circular distance in Figure 3-3C) recall error was systematically shifted clockwise (negative), while counterclockwise distracters led to a

counterclockwise shift in recall error. Indeed, there was a significant interaction between visibility and target-distracter similarity ($F(7,161) = 7.15$, $p < .01$, $\eta^2 = .166$) reflecting that there was a significant effect of target-distracter similarity in the supraliminal ($F(7,161) = 8.18$, $p = .001$, $\eta^2 = 1.836$) but not the subliminal ($F(7,161) = 1.91$, $p = .104$) distracter condition.

To further characterize the shift in recall error I fit a sine wave[6] to the group mean of recall error to get an estimate of the two peak x and y coordinates, i.e. the clockwise and counterclockwise target-distracter difference at which the shift in recall error toward the distracter value is greatest. This revealed that the bias effect was largest following distracters presented 46° (45°) clockwise (counterclockwise) to the target item, with responses shifted 4° toward the distracter value (see Figure 3-5).

Analysis of performance in the identification task revealed that one participant was able to discriminate between stimulus present and absent trials, and their individual subjective ratings mirrored this pattern of performance. Effects were the same when I re-ran the analysis excluding this one participant.

[6] $y = P1*(\sin(2* \pi*x/P2 + 2* \pi/P3)) + P4$, where P1 is the sine wave amplitude (in units of y), P2 is the sine wave period (in units of x), P3 is the sine wave phase (in units of x), and P4 is the sine wave offset (in units of y).

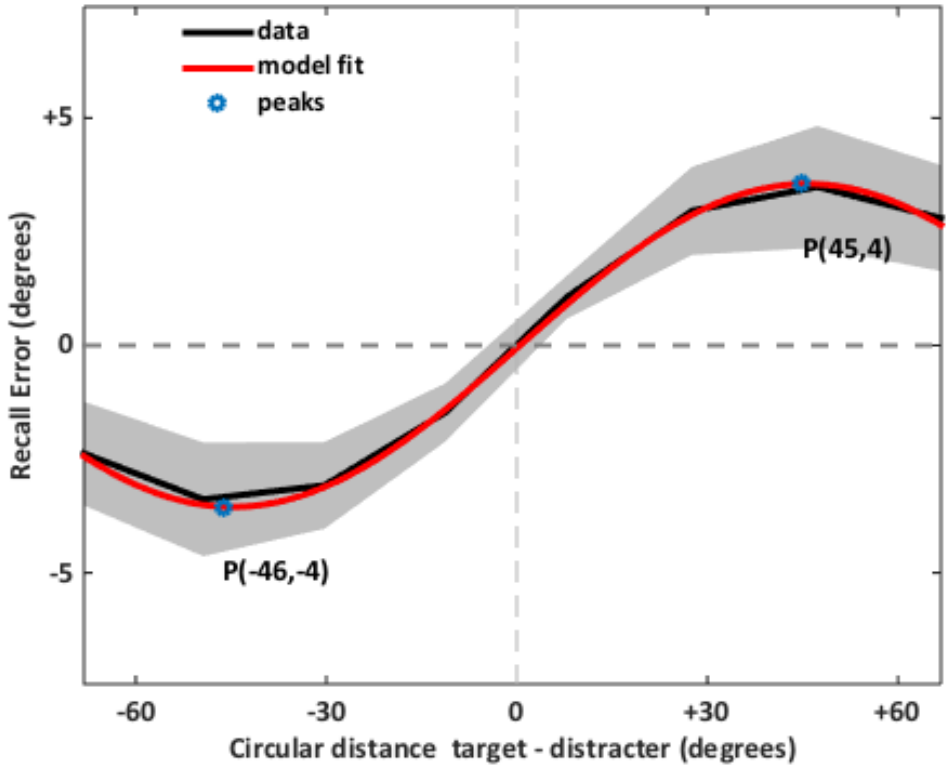

Figure 3-5. Sine wave model fit to the group mean recall error as a function of target-distracter similarity in Experiment 2. Recall error was greatest for distracters presented 46° (45°) clockwise (counterclockwise) to the target item, with responses shifted 4° toward the distracter value.

Control analyses

Previous studies using similar WM tasks have shown that a non-trivial proportion of overall error rates can be attributed to participants mistakenly reporting the distracter item, i.e. misbinding of distracter and target identities (Bays et al., 2009). To ensure that the systematic shift in recall error towards the distracter orientation was not exclusively driven by these misbinding trials, I used a mixture-model approach to model different sources of error contributing to the overall distribution of responses (see Bays et al., 2009 for details) using publicly available Matlab code (bayslab.com). Mixture modelling was conducted in Matlab, using Matlab's

Statistics Toolbox (MATLAB and Statistics Toolbox Release 2012a, The MathWorks, Inc., Natick, MA, 2000) and the Circular Statistics Toolbox for Matlab (Berens, 2009). First, I modelled different sources of error contributing to the overall distribution of responses. The model attributes the overall distribution of responses to a mixture of three separate sources: probability of reporting the target orientation, reporting the distracter orientation, or responding at random. Orientations are recalled with Gaussian variability. This model can be described by the following equation:

$$p(\hat{\theta}) = a\Phi_K(\hat{\theta} - \theta) + \beta \frac{1}{m}\sum_{i}^{m} \Phi_k(\hat{\theta} - \varphi_i) + \gamma \frac{1}{2\pi}$$

where θ is the actual orientation of the target, $\hat{\theta}$ is the reported orientation, Φ_K is the von Mises distribution (circular analogue of the Gaussian) with mean zero and concentration parameter k. The probability of reporting the target is given by a, the probability of reporting the distracter is given by β, and $\varphi_1, \varphi_2, \ldots, \varphi_m$ are the orientations of the m non-target items. The probability of randomly responding is given by $\gamma = 1 - a - \beta$. The von Mises distribution for angle x is described by:

$$f(x|\mu, k) = \frac{e^{k\cos(x-\mu)}}{2\pi I_0(k)}$$

where $I_0(k)$ is the modified Bessel function of order 0, μ is the mean and k is kappa (μ and 1/k are analogous to mean and variance in the normal distribution, respectively). Maximum-likelihood estimates of each parameter (a,β,γ, and k) were obtained by using an expectation-maximization algorithm. This procedure was repeated multiple times using various initial parameter values to ensure the global maximum for each parameter was

found. The original code returns parameter estimates for the probability of reporting the target (a), the probability of reporting the distracter (β), and the probability of responding randomly (γ), on the basis of the entire dataset. However, to estimate these probabilities the relative weights for target, distracter, and random responses are calculated for each trial separately. I used these trial-wise weights to identify target responses, misbinding trials, and guessing trials. Specifically, trials with larger values for distracter weights, i.e. misbinding trials, and guess trials, were excluded from analysis. All misbinding trials had distracter weights of > 0.90 while all target-related trials had distracter weights of < .10. When I re-run the main analyses (calculating recall precision and recall error as before) I observed the same effects as reported above, but with a reduction in overall recall error (significant effect of target-distracter similarity in the supraliminal ($F(7,161) = 3.24$, $p = .009$, $\eta^2 = 0.864$) but not the subliminal ($F(7,161) = 1.48$, $p = .211$) distracter condition. Therefore, the systematic shift in recall error towards the distracter orientation was not due to distracter intrusions, but due to a bias in target-related responses.

Identification Task

Accuracy and d'

Figure 3-4B shows accuracy and d' in the forced-choice discrimination task for the stimulus-present condition as a function of task session. Overall accuracy in the forced-choice discrimination task in the stimulus present condition was not significantly different from chance level (50%): $t(23) = 1.60$, ns; M = 51% ± 8%). A repeated-measures ANOVA with the factor "task session" ('1', '2', '3', '4') revealed that accuracy did not vary

with the task session ($F(3,69) = .243$, ns). Similarly, d' in the stimulus-present condition was not significantly different from zero ($M = .04 \pm .214$, $t(23) = .425$, ns) and showed no effect of task session ($F(3,69) = .048$, ns). These results suggest that, at the group level, participants were not better than chance at performing the discrimination task for the masked stimulus. At the individual subject level one participant performed significantly better than chance (significant binomial test on accuracy data at the .05 level [$p = .044$]), suggesting they could differentiate between present and absent trials to some extent.

Subjective Ratings

A Wilcoxon signed-ranks test comparing average awareness ratings in the stimulus-present and stimulus-absent condition indicated that ratings were higher in the present condition ($M = 1.368 \pm 0.08$) than in the absent condition ($M = 1.316 \pm 0.07$), $Z = -2.40$, $p = .016$, see Figure 3-4C).

To examine further what was driving these differences, I calculated the proportion of each rating for each participant separately, and then ran four separate Wilcoxon signed-ranks tests comparing the proportions of ratings between the present and absent condition for each rating option. "Not Seen" ratings occurred significantly more often in the absent condition ($M = 0.73 \pm 0.06$) than in the present condition ($M = 0.68 \pm 0.06$), $Z = -2.63$, $p = .01$). Conversely, proportions of "Weak Glimpse" ratings were significantly higher in the present condition ($M = 0.268 \pm 0.05$) than in the absent condition ($M = 0.222 \pm 0.05$, $Z = -2.64$, $p = .008$). Proportions of "Almost Clear" and "Absolutely Clear" did not differ between the absent condition ($M = 0.039 \pm 0.02$, and $M = 0.01 \pm 0.004$, respectively) and the

present condition (M = 0.04 ± 0.02, and M = 0.01 ± 0.003, respectively), all p's > 0.5. Figure 3-4D shows the proportion of awareness ratings for the stimulus-present and stimulus-absent conditions separately.

Together, these results suggest that, at the group level, participants were unable to discriminate stimulus orientations on present trials, reinforcing the results obtained in the d' and accuracy analyses. However, analysis of the subjective ratings suggests that participants could differentiate between stimulus-present and stimulus-absent trials to a small degree, since participants reported not seeing a stimulus significantly more often when there was no stimulus compared to when there was. Conversely, when a stimulus was present, participants reported perceiving a weak glimpse more often than when there was no stimulus.

Importantly, subjective ratings converged with objective performance in the forced-choice identification task where participants did not perform better than chance. The only significant differences between stimulus present and absent trials in subjective ratings was between "Not Seen" and "Weak Glimpse" ratings, which participants were instructed to use to describe perceptual experiences that were lacking any orientation information. Conversely, ratings of "Almost Clear" and "Absolutely Clear", which participants had been instructed to use when they had experienced orientation information, did not differ between conditions. Thus, subjective ratings indicated that participants could detect, but not identify, stimulus presence, which is consistent with the forced-choice identification task measures.

3.6.4 Interim Discussion

The results show that supraliminal but not subliminal distracters presented before the target orientation strongly influenced the observers' subsequent recall based on the WM representation. When distracters were supraliminal, observers' reports of the target orientation were strongly biased towards the distracter orientation, and the strength of this bias varied with the similarity between distracter and target orientations. The effect represented a systematic bias rather than intrusions from trials on which participants incorrectly reported the orientation of the distracter stimuli (misbinding errors). This occurred despite the fact that the distracter was task-irrelevant and never probed. When presented subliminally, distracters exerted no influence on orientation reports.

The results, therefore, suggest that subliminal stimuli are unable to influence WM representations that become available for conscious report. However, the fact that the masked stimuli in the current task were consistently task-irrelevant and could be completely ignored may have dampened any influence they may have otherwise exerted. It is well known that task relevance can modulate stimulus processing (Ansorge & Neumann, 2005; Gayet et al., 2013b; Gorgoraptis et al., 2011; Wyart, Nobre, & Summerfield, 2012; Zokaei, Manohar, Husain, & Feredoes, 2014). In particular, task relevance has been shown to mediate the effect of subliminal stimuli in some tasks (Ansorge & Neumann, 2005; Gayet et al., 2013b). Gayet et al. (2013b) asked participants to complete a peripheral target-detection task with central arrow cues that were either subliminal or supraliminal (randomly intermixed), with the majority of cues being

supraliminal. Subliminal cues were never predictive, while supraliminal cues could be non-predictive or highly predictive (presented as separate experimental conditions on separate days). It was found that subliminal arrow cues only facilitated performance when they were presented among predictive supraliminal cues. This facilitatory effect increased over the course of the experiment, suggesting that the usage of subliminal cues was based on participants learning the predictive value of the supraliminal cue. Subliminal information only appears to be used when it is presented in a context where there is a reason to use it.

Thus, in the next experiment, I changed the task so that the subliminal orientation stimulus was task-relevant. Specifically, participants now had to report the orientation of the masked item on a proportion of trials even when it was presented subliminally. I hypothesized that task relevance would encourage processing of the masked subliminal stimuli, thus allowing them to bias orientation reports on subsequent targets.

3.7 Experiment 3: Influence of task-relevant subliminal vs supraliminal stimuli on WM

3.7.1 Methods

Participants

Twenty-one new subjects took part in the experiment (10 male, 18 – 32 years, two left-handed). Two additional participants were tested but their data were not included in the analysis as their task performance was very poor (high error rates and >50% misbinding or guess trials).

Stimuli

The stimuli were identical to those used in Experiments 1 and 2.

Task

As in Experiment 2, participants performed two tasks: A WM task and an identification task. The WM task (Figure 3-6A) was adapted from the task used in Experiment 2. The main difference was that participants were prompted to report the orientation of either the masked or unmasked stimulus. This was done to render the masked stimulus task-relevant, and allowed me to assess whether the orientation of the task-relevant subliminal stimulus could influence encoding of the unmasked item. Concurrently, this manipulation also allowed me to assess whether participants had any information available about the subliminal orientation when asked to reproduce it. The purpose of the identification task was to ensure that the pattern mask was effective at rendering stimuli subliminal.

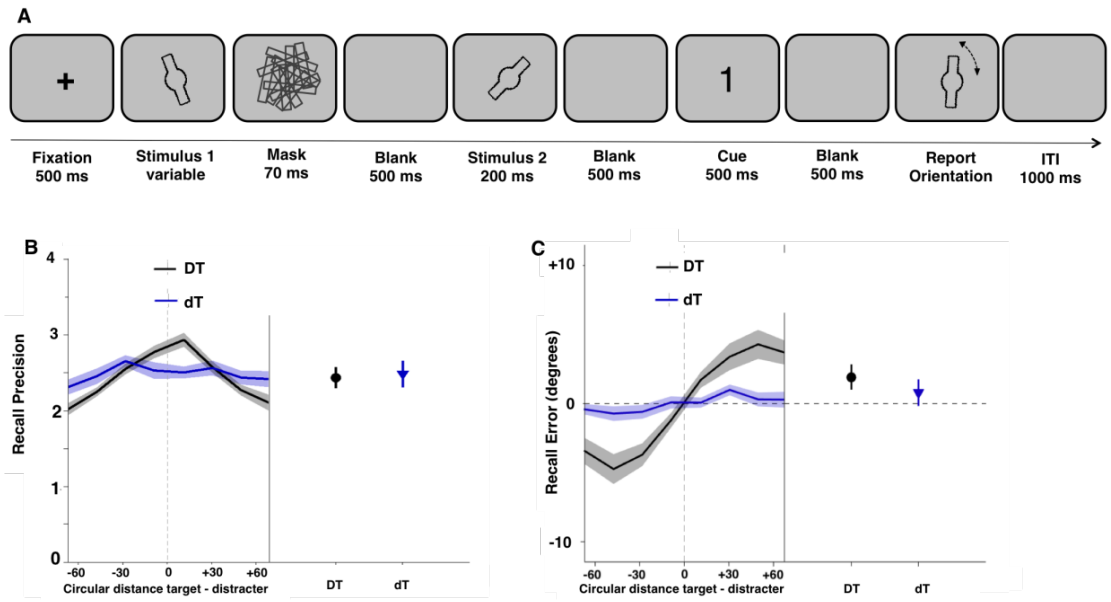

Figure 3-6. Design and results of the WM task used in Experiment 3. **A**, A central fixation cross was presented for 500 ms followed by a first orientation stimulus, followed by a 70-ms mask. The presentation duration of the first orientation stimulus was varied to yield two conditions of equal probability: one in which the orientation was visible despite the mask and another in which the mask rendered the orientation stimulus invisible. After

a 500-ms delay, the second orientation stimulus was presented for 200 m (unmasked). This was followed by another 500-ms delay, and then by a 500-ms cue stimulus, which indicated to participants which orientation they had to report at the end of the trial. After a final 500-ms delay, participants reported the cued, target orientation by matching the orientation of a probe stimulus with the cued orientation stimulus using the computer mouse. **B**, Mean recall precision of cued orientation recall as a function of cued and uncued orientation similarity (left) and mean overall recall precision of cued orientation recall in experiment 3 plotted for both conditions separately (N = 19). **C**, Mean recall error as a function of cued and uncued orientation similarity in experiment 3 plotted for the DT and dT condition separately (left), and mean overall recall error for the two conditions separately (N = 19).

Each trial began with the presentation of a central fixation cross for 500 ms followed by a first orientation stimulus, which was followed by a 70-ms mask. The presentation duration of the first orientation stimulus was varied to yield two conditions of equal probability: a first condition in which the orientation was visible despite the mask and another in which the mask rendered the orientation stimulus invisible (duration individually determined, M = 32ms ± 6). After a 500-ms delay, the second orientation stimulus was presented for 200 ms (unmasked). This was followed by another 500-ms delay, and then by a 500-ms cue stimulus, which indicated to participants which orientation they had to report at the end of the trial. The cue was the number ‘1’ or ‘2’, denoting the first or second orientation stimulus, respectively. After a final 500-ms delay, participants reported the cued, target orientation by matching the orientation of a probe stimulus with the orientation of the cued stimulus using the computer mouse. Participants also completed an identification task that was identical to the one used in Experiment 2 (see Figure 3-4A).

Design and Procedure

Design and procedure were identical to Experiment 2 apart from the following changes. The WM task consisted of a fully factorial design with two variables: stimulus 1 visibility (subliminal, supraliminal) and cue (stimulus 1, stimulus 2). The different trial types were randomly intermixed within a block. Thus, the experimental design yielded four conditions: supraliminal distracter presented before a supraliminal target (I will refer to this condition as DT), subliminal distracter presented before a supraliminal target (dT), supraliminal target presented before presentation of a supraliminal distracter (TD), and subliminal target presented before a supraliminal distracter (tD).

Before starting the main experimental task, participants completed 50 practice trials in which they received feedback after every trial. Following practice, participants completed 24 blocks of 50 trials (1200 trials in total) of the WM task, yielding 300 trials in each condition. The first 600 trials were completed in the first session and the last 600 trials in the second session. The identification task was identical to that in Experiment 2.

Staircase Procedure

The staircase procedure was the same as the one used in Experiment 2.

3.7.2 Analysis

Working Memory Task

Effects on recall precision and recall error were tested using ANOVAs and t tests. First I tested whether task-relevant subliminal and supraliminal stimuli presented before the target can influence orientation reports. To that end I ran a 2-by-8 repeated-measures ANOVA with the

variables “visibility” (subliminal, supraliminal) and “target-distracter similarity” (8 bins ranging from 90° clockwise from the target to 90° counter-clockwise) on recall precision and recall error. In a second analysis, I examined whether a supraliminal distracter presented after the target can influence reports on the target using a one-way repeated-measures ANOVA with the variable “target-distracter similarity” on recall precision and recall error. Finally, I tested whether participants had any information available when asked to report a subliminal target presented before a supraliminal distracter. I reasoned that if participants had no orientation information available regarding the subliminal stimulus, the mean error should be high and not different from chance performance. I did not analyse recall precision, as in this case this measure might be contaminated: if a participant consistently reported the uncued, supraliminal orientation, recall precision would be high as errors in orientation report cluster around a certain value (the uncued orientation) and error variability is low in each bin. Recall error, on the other hand, would still be high and is therefore more informative about task performance.

Identification Task

Performance in the identification task was assessed as described in Experiment 2.

3.7.3 Results

Effects of subliminal and supraliminal distracters presented before targets

Recall precision

Figure 3-6B shows the precision with which participants recalled the target orientation as a function of target-distracter similarity for the dT and the DT conditions.

There was a strong trend for a significant difference between dT and DT condition in overall precision (M_{dT} = 2.31 radian-1 ± 0.171 and M_{DT}= 2.12 radian-1 ± 0.137. t(20) = -1.97, p = .064, d = 0.40). Similarly, as in Experiment 2, the interaction between visibility and target-distracter similarity was significant (F(7,140) = 4.18, p < .001, η^2= 0.173). To examine this interaction further, two separate one-way repeated-measures ANOVAs assessed the effect of target-distracter similarity for each condition separately. These revealed a significant effect of target-distracter similarity on recall precision in the supraliminal (DT) condition (F(7,140) = 9.09, p < .001, η^2= .313) but not in the subliminal (dT) condition (F(7,140) = 1.13, p = .347). Specifically, the more similar the orientation of the supraliminal distracter was to the target orientation, the better was participants' recall precision (see Figure 3-6B).

The main effect of target-distracter similarity on precision was also significant (F(7,140) = 5.61, p < .001, η^2= 0.219), but there was no main effect of visibility (F(1,20) =.56, ns).

Recall error

Figure 3-6C shows recall error for the target orientation as a function of target-distracter similarity for the dT and DT condition separately. As in the previous experiment, there was a significant interaction between visibility by target-distracter similarity (F(7,140) = 5.49, p = .001, η^2= 1.507), reflecting a significant effect of target-distracter similarity in the supraliminal

(DT) ($F(7,140) = 10.50$, $p < .001$, $\eta^2 = 2.41$) but not the subliminal (dT) ($F(7,140) = 1.21$, $p = .310$) condition. When the supraliminal distracter orientation was counter-clockwise to the target orientation (a negative circular distance in Figure 3-6C) there was a counter-clockwise (negative) shift in participants' orientation reports of the target. Similarly, when the supraliminal distracter orientation was clockwise to the target orientation (a positive circular distance in Figure 3-6C) there was a clockwise (positive) shift in participants' orientation reports of the target. Subliminal distracters did not influence WM recall.

To further characterize the shift in recall error I fit a sine wave to the group mean of recall error in the DT condition, as in Experiment 2. This revealed that the bias effect was largest following distracters presented 49° (48°) clockwise (counterclockwise) to the target item, with responses shifted 4° toward the distracter value (see Figure 3-7).

Analysis of performance in the identification task revealed that one participant was able to discriminate between stimulus present and absent trials, with the individual subjective ratings mirroring this pattern of performance. I re-ran the analysis excluding this participant (N = 20) and the results remained equivalent. Thus, I replicated the findings of the previous experiment that participants' orientation reports are systematically shifted towards the distracter orientation if and only if the distracter is visible.

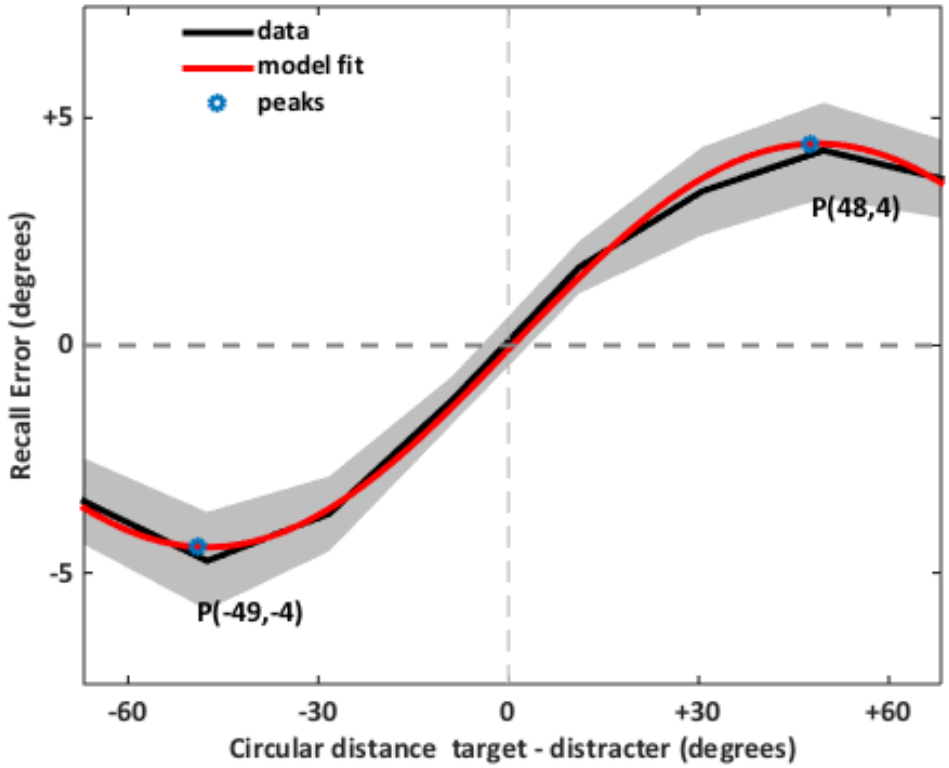

Figure 3-7. Sine wave model fit to the group mean recall error as a function of target-distracter similarity in condition DT in Experiment 3. Recall error was greatest for distracters presented 49° (48°) clockwise (counterclockwise) to the target item, with responses shifted 4° toward the distracter value.

Control analyses

To ensure that the systematic shifts in recall error with distracter orientation observed in Experiment 3 were not exclusively driven by trials in which participants incorrectly responded based on distracter stimuli (misbinding trials), I applied mixture-modelling to the data set, and excluded misbinding and guess trials. I then re-run the main analyses on recall precision and recall error. The same effects were observed (significant effect of target-distracter similarity in condition DT but not condition dT: $F(7,140)= 8.60$, $p < .001$, $\eta^2= 2.105$; and $F(7,140) = 0.85$, $p = .518$, respectively).

Effects of subliminal and supraliminal distracters presented after targets

Recall precision

Figure 3-8A shows the precision with which participants recalled the target orientation as a function of target-distracter similarity for the TD condition. There was no significant effect of target-distracter similarity on recall precision in the TD condition ($F(7,140) = 1.87$, $p = .210$, $\eta^2 = .11$).

Recall error

Figure 3-8C shows recall error for the target orientation as a function of target-distracter similarity for the TD condition. Visual inspection of the graphs suggests that recall error varied as a function of target – distracter difference. As before, when the distracter orientation was to the right of the target orientation there was a rightwards shift in participant's orientation reports of the target, and vice versa. Statistical analysis of the data supported this and there was a strong trend for an effect of target-distracter similarity ($F(7,140) = 2.60$, $p = .057$, $\eta^2 = .805$). This suggests that presentation of a supraliminal distracter can influence the report of a remembered orientation, irrespective of whether it is presented before or after target orientation. The biasing effect in the TD condition was equivalent to the biasing effect in the DT condition ($t(20) = -1.00$, $p = .33$). To further characterize the biasing effect I fit a sine wave to the group mean of recall error in the TD condition. This revealed that the bias effect was largest following distracters presented 36° (45°) clockwise (counterclockwise) to the target item, with responses shifted 4° toward the distracter value (see Figure 3-9).

Figure 3-8B shows mean recall error in the tD condition for each subject separately. The mean error in the tD condition was 89° (±0.9°), which was not significantly different from chance (i.e. 90°; t(20) = -1.29, p = .209). This corroborates the prime identification results and shows that stimuli remain subliminal in the actual task, where reports of subliminal stimuli are randomised.

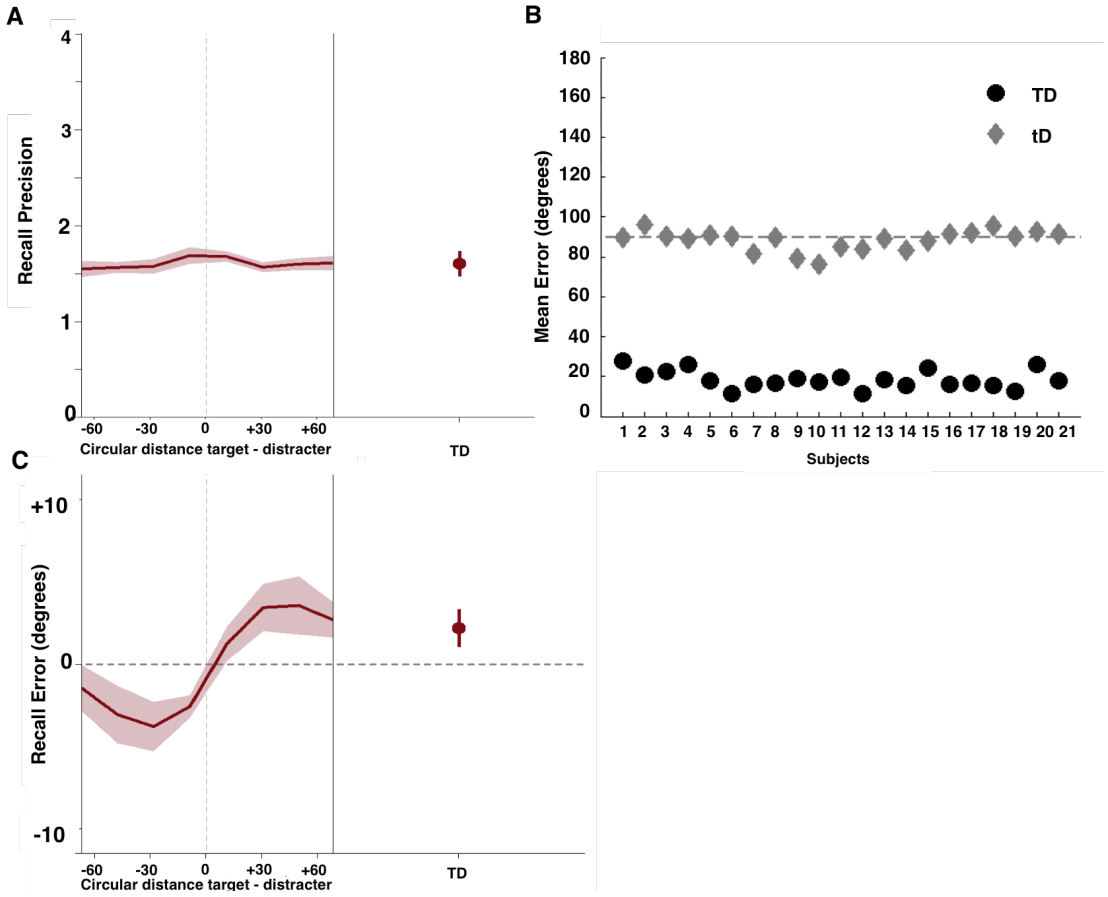

Figure 3-8. A, Mean recall precision of cued orientation recall as a function of cued and uncued orientation similarity in Experiment 3 for the TD condition (left), and average recall precision of cued orientation recall for the same condition (N = 19). **B**, Mean error in reporting the cued orientation for the TD and tD conditions in Experiment 3 plotted for each subject separately. The dotted grey line indicates chance performance. **C**, Mean recall error as a function of cued and uncued orientation similarity for the TD condition (left), and mean overall recall error for the two conditions separately (N = 19).

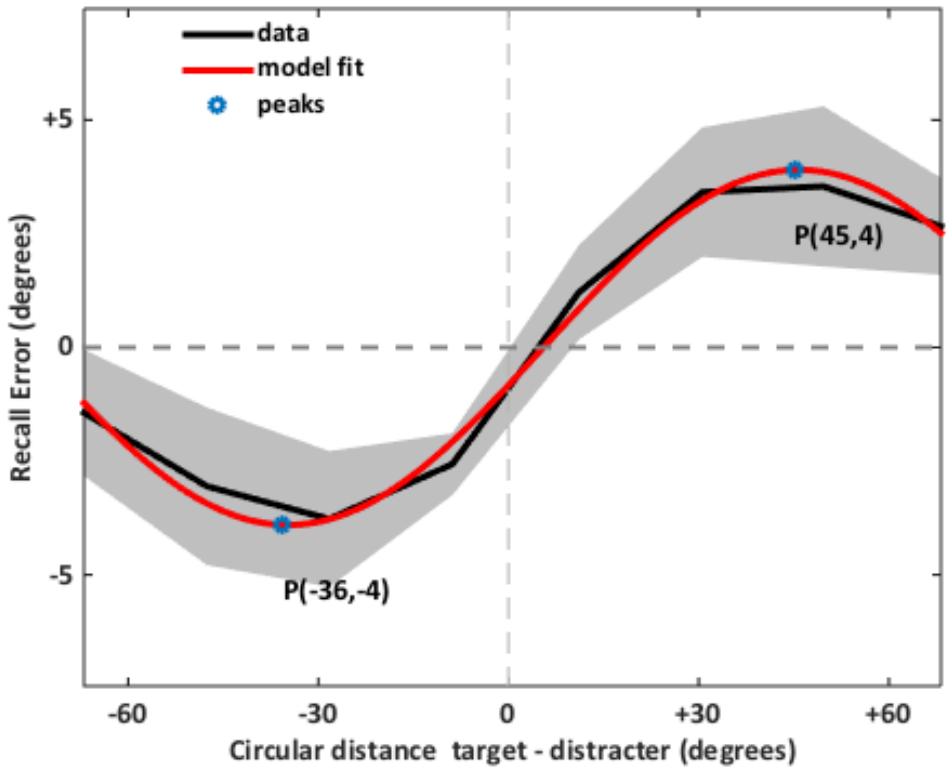

Figure 3-9. Sine wave model fit to the group mean recall error as a function of target-distracter similarity in condition TD in Experiment 3. Recall error was greatest for distracters presented 36° (45°) clockwise (counterclockwise) to the target item, with responses shifted 4° toward the distracter value.

Control analysis

To ensure that the systematic shifts in recall error with distracter orientation observed in Experiment 3 were not exclusively driven by trials in which participants incorrectly responded based on distracter stimuli (misbinding trials), I applied mixture-modelling to my data set, and excluded misbinding and guess trials. I then re-run the main analyses on recall precision and recall error. The qualitative pattern of results remained the same but the trend for a significant effect of target-distracter similarity in condition TD was now significant: $F(7,140) = 3.02$, $p = .026$, $\eta^2 = 0.918$.

Individual differences in recall error performance

At the group level, distracter presentation always biased target recall toward the distracter value in Experiment 2 and both conditions in Experiment 3. However, there were pronounced individual differences not only in the magnitude but also in the direction of the effect (see Figure 3-10). In both Experiments, some participants exhibited the opposite effect and target recall was biased away from the distracter value.

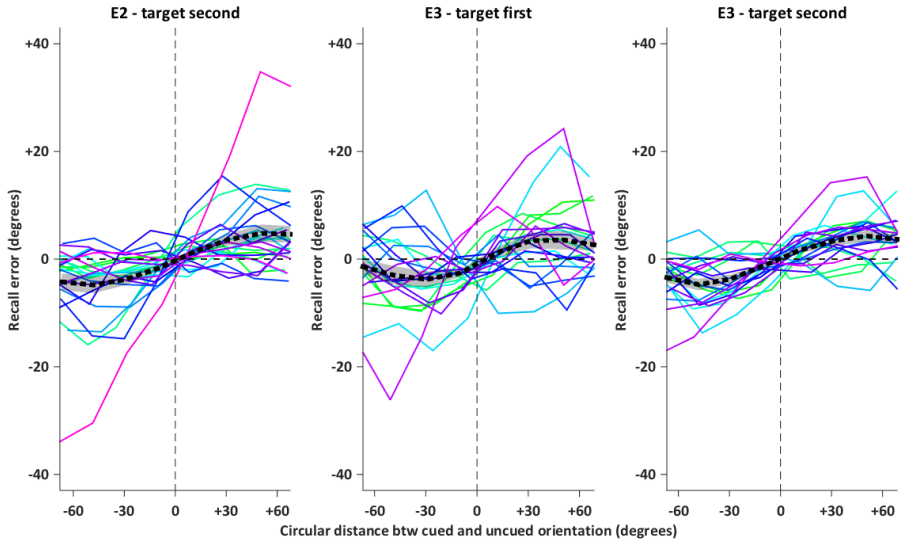

Figure 3-10. Individual participant's recall error performance in Experiments 2 and 3. Individual performance is shown in the coloured lines, while the group mean (± 1 standard error of the mean) is shown as the black dotted line (grey shading). Recall error was significantly biased toward the distracter value at the group level, but in each Experiment there was a subset of participants that displayed the opposite pattern.

Identification Task

Accuracy responses were not recorded during two of the four sessions for one participant due to technical errors and only the subjective ratings were included in the analysis. For another participant, both

subjective ratings and accuracy responses from one session were not included in the analysis as they were accidentally overwritten.

Accuracy and d'

Figure 3-11A shows accuracy and d' in the forced-choice discrimination task for the stimulus-present condition as a function of task session. Accuracy in the stimulus present condition was not significantly different from chance level (50%): $t(18) = .35$, ns; M = 50% ± 1%), and did not vary with the task session ($F(3,54) = .73$, ns). Similarly, d' in the stimulus present condition was not significantly different from zero ($t(18) = .37$, ns; M = .02 ± .06). Performance did not vary over task blocks/sessions ($F(3,54) = .96$, ns). These results suggest that, at the group level, participants were not better than chance at performing the discrimination task for the masked stimulus. At the individual subject level, one participant performed significantly better than chance (significant binomial test on accuracy data at the .05 level [$p < .001$]), suggesting they could differentiate between present and absent trials to some extent.

Subjective Ratings

Figure 3-11B shows the average awareness ratings for the stimulus present and absent conditions for each testing session. A Wilcoxon signed-ranks test comparing average awareness ratings in the stimulus-present and stimulus-absent condition indicated that overall ratings were not different between the present and absent condition (M = 1.45 ± 0.10; M = 1.48 ± 0.10, respectively). Furthermore, this did not change over the course of the experiment, and there was no difference between conditions in any of the sessions (all Z's < 1).

Since there were no significant differences in average awareness ratings between stimulus-present and stimulus-absent conditions I did not statistically evaluate the proportions of different ratings given in each of the two conditions. However, for completeness the proportions of ratings are depicted in Figure 3-11C.

Together these results reinforce the results in the accuracy and d' analyses, suggesting that participants were unable to discriminate stimulus orientations on present trials at the group level. Furthermore, analysis of the subjective ratings suggests that participants could not differentiate reliably between stimulus-present and stimulus-absent trials as participants.

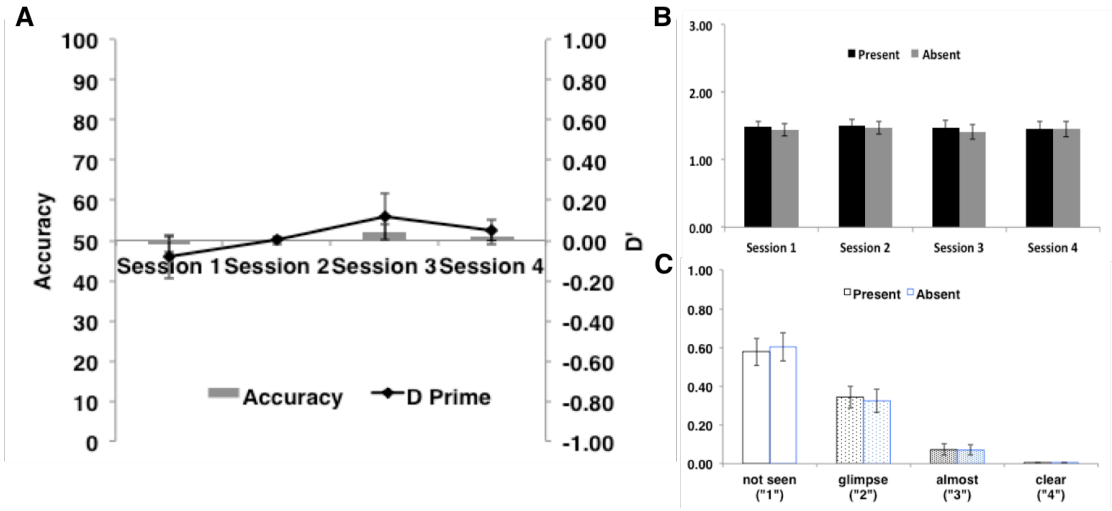

Figure 3-11. A, Accuracy and d' in the forced-choiced identification task for stimulus-present condition as a function of session number in Experiment 3. Error bars reflect ± 1 standard error of the mean. **B**, Mean awareness ratings for the stimulus present and stimulus absent condition in Experiment 3 for each session separately. **C**, Proportion of awareness ratings in the identification task for stimulus-present and stimulus-absent condition.

3.7.4 Interim Discussion

Experiment 3 replicated and extended the findings of Experiment 2: a distracter orientation influenced observers' report of a remembered target orientation, whether it was presented before or after the target orientation. Observers' reports of the target orientation were biased towards the distracter orientation. In contrast, the orientation of subliminal distracters did not influence responses. This was the case even when the distracter orientation was relevant for task completion – previously suggested to be a necessary condition for subliminal information to influence subsequent perception (Gayet et al., 2013b).

3.8 Discussion

I have shown, using a novel and sensitive approach to measure influence of irrelevant stimuli on targets, that irrelevant distracters can influence WM performance to target stimuli when distracters are supraliminal but not when they are subliminal according to objective visibility measures. Experiments 1a and 1b replicated effects of subliminal stimuli on perceptual judgements (Eimer & Schlaghecken, 1998; 2002), while Experiment 2 demonstrated that these effects do not extend to WM decisions. Experiment 3 replicated the effects in Experiment 2 even when the subliminal stimuli were made task relevant. Interestingly, it also revealed that a supraliminal distracter orientation influences memory recall for a target orientation both when it appears before and when it appears after the target orientation. Whether this effect represents an encoding, response or maintenance effect is difficult to determine with the current set of results, and I leave this open to interpretation. In addition, the presentation of

supraliminal distracters also affected WM recall precision and task performance was worse when distracters were presented relative to when distracters were absent, or presented subliminally.

The present finding on subliminal stimuli and WM contrasts with findings by Silvanto and Soto (2012), who reported an effect of a subliminal distracter on WM representations of a target item. Participants retained a target orientation in memory to compare against a probe orientation stimulus. During the delay period, a masked subliminal distracter was presented on most trials, which was either congruent or incongruent with the target orientation. In the incongruent condition, but not the congruent or no-distracter conditions, participants' accuracy was significantly impaired. A critical factor here, however, may be the level of awareness that participants had on the masked distracter. Crucially, Silvanto and Soto (2012) defined levels of awareness based on subjective measures; a stimulus was classified as subliminal when observers reported not having seen it. I, on the other hand, classified stimuli as subliminal by presenting them in such a way that observers could not identify them above chance. Thus, my study defined stimuli as subliminal on the basis of an objective measure, while Silvanto and Soto (2012) used a subjective measure. The difference in results suggests that information of which observers are subjectively not aware can influence WM representations as long as observers show some sensitivity ($d' > 0$) to the stimulus. However, there is no evidence of infiltration of WM for stimuli that are objectively below a threshold for awareness. The current results are consistent with the notion that items should be available for awareness to be encoded, or to influence, other WM

representations (cf. Baddeley, 1986; Bundesen, 1990). The results may further suggest that WM is not simply a matter of activating sensory codes (cf. Cowan, 1995) regardless of their level of awareness.

It is important to remember that the notion of subliminal processing remains controversial (Pratte & Rouder, 2009; Reingold, 2004) The two replications of the influence of subliminal primes on speeded responses in a visuomotor task similar to that used by Eimer and Schlaghecken (1998, 2002) add support to the possibility of subliminal effects, at least in the context of influencing motor tendencies in perceptual tasks. The aim, in Experiments 2 and 3, was to test for putative effects of subliminal stimuli in the context of WM, but these experiments also differed from those in Experiment 1 in terms of the types of representations required to guide responses.

Indeed, an alternative explanation of the present set of results might focus less on the distinction between WM and perception, but instead on the nature of representations that observers need to access to complete the task. Notably, in Experiments 1a and 1a I implemented a strict response deadline and only stimulus-response assignments needed to be accessed. In Experiments 2 and 3, on the other hand, the perceptual content of the stimulus needed to be retrieved as observers had to report the visual appearance of a stimulus. Thus, it might be that subliminal effects are not readily observed in WM-type tasks because decisions require consideration of the appearance of an item, whereas observers' responses can be modulated by subliminal information in tasks that require access to stimulus-response assignments. Future research could address this possibility by

measuring subliminal effects in perceptual tasks that require access to the perceptual content of the stimuli, and subliminal effects in WM tasks that do not require access to it.

Mutual influences between supraliminal stimuli on observers' reports have recently been reported in tasks using spatial frequency (Dubé et al., 2014; Huang, 2010) and orientation stimuli (Fischer & Whitney, 2014). Huang and Sekuler (2010) asked participants to reproduce the spatial frequency of one of two successive Gabors, as indicated by a cue stimulus presented after the Gabors. The reproduced spatial frequency of the target item was influenced by the spatial frequency of the non-target stimulus, and responses were biased in the direction of the non-target item. This effect was found both within a trial, between a target and distracter stimulus, and between trials, from one target item to the subsequent target item. Similarly, Fischer and Whitney (2014) found a systematic bias in observers' perceived orientations in the direction of previously seen orientations. Thus, this biasing effect appears to be robust across different experimental set ups and stimuli, and seems to be a general principle of information processing.

Importantly, here I explicitly controlled for the possibility that responses incorrectly based on the wrong stimulus may contaminate results and drive such biasing effects. The studies reviewed above did not model 'misbinding'. What appears to be a biasing effect in the direction of the distracter might emerge when observers mistakenly report the distracter instead of the target on some small number of trials. My analyses modelling and excluding misbinding and guessing trials unambiguously confirm a real influence of distracter stimuli on the memory for the target.

Unlike in perceptual and visuomotor processes, these WM biases were not induced by subliminal information. Gayet et al. (2013b) suggested that subliminal information is only used when the context in which it is presented provides an incentive to make use of this information. However, I did not observe an effect of subliminal distracters on target recall even when observers were asked to report back the distracter's orientation on some trials. Instead, I argue that WM representations are not as easily influenced by subliminal material and task-relevance may have a stronger mediating effect on the processing of subliminal stimuli in perceptual tasks, than in WM tasks such as ours.

I show an effect of a distracter on memory representations for a target item over and above misbinding in two separate experiments. This biasing effect was observed even though only two items (the target and the distracter) were presented, which is well below capacity limitations of WM. This finding suggests a permeability of memory representations even before capacity is exceeded, constraining current models of WM.

4 Effects of Retrospective Attention and Contrast on Working Memory Performance

4.1 Chapter Abstract

Orienting attention retrospectively to selective contents in WM influences performance. A separate line of research has shown that stimulus strength shapes perceptual representations. There is little research on how stimulus strength during encoding shapes WM performance and interactions between WM items, and how effects of retrospective attentional orienting might vary with changes in stimulus strength. In this chapter, I explore these questions in three experiments using a continuous-recall WM task. In Experiment 1, I show that benefits of cueing spatial attention retrospectively during WM maintenance (retrocueing) varies according to stimulus contrast during encoding. Retrocueing effects on overall recall performance emerge for supraliminal but not subliminal stimuli. However, once stimuli are supraliminal performance is no longer influenced by stimulus contrast. In Experiments 2 and 3, I used a mixture-model approach to examine how different sources of error in WM are affected by stimulus contrast and retrocueing, and how interactions among WM items change as stimulus contrast during encoding changes. For high-contrast stimuli (Experiment 2) retrocues increased the representational quality of successfully remembered target items. For low-contrast stimuli (Experiment 3) retrocues decreased the probability of mistaking a target with distracters. Across all experiments, effects were modulated and predicted by individual differences in baseline performance. These results suggest that the processes by which retrospective attentional orienting shape WM

performance are flexibly adapted to the quality of WM representations, which in turn depends on stimulus strength and individual differences in WM capacity.

4.2 Introduction

WM – our ability to maintain information over short periods of time to guide adaptive behaviour – is a fundamental cognitive function that is constrained by its limited capacity, shaping both the quality and quantity of WM contents. Given these capacity restrictions, researchers have been concerned with identifying and characterising variables that determine WM variability within individuals over time.

For one, there are systematic effects of orienting attention retrospectively to individual locations or items within WM: an effect termed 'retrocueing' benefit (Griffin & Nobre, 2003; Landman et al., 2003; Makovski et al., 2008; Matsukura et al., 2007; Sligte, Scholte, & Lamme, 2008). Studies that addressed functional mechanisms underlying the retrocueing benefit have produced mixed results. To date it is unclear how and when retrocues reduce which source of error in WM. In the General Introduction, I suggested this variability in reported retrocueing benefits is related to variability in WM. For example, it has been proposed that intrinsic noise fluctuations impact on the stability of WM representations (Bays, 2014; Fougnie, Suchow, & Alvarez, 2012a; Matthey et al., 2015; van den Berg et al., 2012). Random fluctuations in internal neural processing influence not only the detection of an external stimulus (Hanslmayr et al., 2007; Mathewson et al., 2009; van Dijk, Schoffelen, Oostenveld, & Jensen, 2008), but also WM recall performance (Myers et al., 2014).

In addition to internal source of noise visual information can also be contaminated by external sources of noise. Fluctuations in bottom-up stimulus properties and viewing conditions (e.g. stimulus contrast or noise levels) during encoding might impact the quality of WM representations. To the best of my knowledge, the influence of stimulus strength during encoding on mental representations has been mostly overlooked in WM research and is typically ignored in most contemporary models of WM (Bays & Husain, 2008; Fougnie, Suchow, & Alvarez, 2012a; Luck & Vogel, 1997; van den Berg et al., 2012; Zhang & Luck, 2008), although it would appear to be an obvious influence on WM variability. However, we know two things: first, basic stimulus properties have been shown to modulate WM recall performance and perceptual judgements. Girshick et al. (2011) reported that orientation judgments are more accurate for cardinal than for oblique orientations. Findings by Bae, Olkkonen, Allred, Wilson, and Flombaum, (2014) show that these effects persist into WM, and recall performance for colours depended on the specific colour value. At least some of the variability in WM appears to be stimulus driven. Second, models of perceptual decision-making explicitly consider the quality of incoming sensory information as a critical variable with direct influence on perceptual representations (e.g. Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Gold & Shadlen, 2007). These findings provide support for the proposal that changes in stimulus strength during encoding may impact on the representational quality of items in WM. Additionally, individual differences both in task-strategy and ability introduce variability to the quality of WM representation across time and between individuals. If effects of retrospective attentional orienting depend on the representational quality of

WM items, there should systematic effects of bottom-up stimulus properties on the retrocueing benefit, and this should be further modulated by individual differences.

Interestingly, a recent computational model proposes that one specific source of WM error, erroneous nontarget responses (i.e. misbinding errors), is related to representational limitations in WM (Matthey et al., 2015). According to the model, different items in WM are represented by one neural population, which leads to overlap between items, which in turn leads to misbinding errors. Importantly, how items overlap is determined during encoding. Therefore, changes in bottom-up stimulus properties should be reflected in changes in misbinding rates. Whether and how changes in the relative proportion of different errors in WM relate to effects of retrospective attentional selection is unclear.

An additional question emerges from the fact that only high-contrast stimuli have been used, which allow for conscious perception of the stimulus. It remains unexplored whether attention could have a 'boosting' effect on subliminal representations when stimuli are presented at a contrast level below the level of conscious awareness. Sergent et al. (2013) asked whether retrospective attention could influence subjective visibility of a low-contrast stimulus that would not be consciously perceived otherwise. Observers were presented with a single target, a Gabor patch of a random orientation, either to the left or right of fixation and with an exogenous cue either before or after presentation of the Gabor. After a short delay, participants reported the target orientation by choosing between two options. Baseline accuracy in orientation discrimination without a cue was set to 80%. Following the orientation discrimination participants rated the

subjective visibility of the target on a scale ranging from “not seen” to “max visibility”. Performance in the orientation discrimination task was significantly improved by a congruent cue both with cues appearing before and after the target. Furthermore, the proportion of trials in which participants reported the target as “not seen” was significantly lower in the congruent compared to the incongruent cueing condition. This, the authors argue, shows that orienting attention to a stimulus that is no longer visible can elicit conscious experience of the stimulus when otherwise the stimulus would have escaped consciousness. While this study provides strong evidence for a boosting effect of retrospective attention on neural activity underlying internal representations, it does not necessarily show that attention can render otherwise invisible stimuli visible. The authors’ conclusion is based on a change in the proportion of subjective ratings of awareness while the target stimulus was staircased to be accurately identified in 80% of trials, well above threshold. Since only subjective measures of awareness were employed there is not sufficient evidence that observers were truly unaware of the target stimulus when they reported zero visibility (Hannula et al., 2005).

In this chapter, I ask how retrocueing combines with the effects of stimulus quality to modulate WM. Can retrocues only rescue items with sufficient contrast to be encoded in the first place, or can they have an effect on WM performance in low-contrast situations, or even when stimuli are presented below a level for conscious awareness? I employed the same task (Figure 4-1) across three experiments in which I varied contrast levels of the to-be-remembered information and manipulated attentional orienting

within WM using a retrocueing paradigm (Griffin & Nobre, 2003; Murray et al., 2013).

4.3 General Methods

4.3.1 Participants

All experimental protocols were reviewed and approved by the Central University Research Ethics Committee of the University of Oxford. All participants had normal or corrected-to-normal vision and were naïve to the purpose of the experiment. Participants gave informed consent before taking part, and received course credits or financial compensation (£10 per hour) for taking part. In Experiment 1, 31 volunteers were included (all right-handed according to self report, mean age 25 years (range 18 – 35), 12 male). Six additional volunteers were tested but excluded from data analysis (see the “Inclusion Criteria” section for details). In Experiment 2, 30 new volunteers were included (all right-handed, mean age 26 years (range 18 – 35), 10 male). Two additional volunteers were excluded from data analysis. In Experiment 3, 34 volunteers were included (all right-handed, mean age 24 years (range 18 – 35), 12 males). Four additional volunteers were tested but excluded from data analysis.

4.3.2 Apparatus

Stimuli were created in MATLAB v.7.10 (MathWorks) and presented using the Psychophysics Toolbox (Kleiner et al., 2007). They appeared against a uniform mid-grey background on a 23-inch light-emitting-diode (LED) display. The display had a spatial resolution of 1920x1080 and a refresh rate of 120 Hz. In Experiment 3, a chin rest was used to maintain a

constant viewing distance and head position, and eye gaze was recorded binocularly with a desktop-mount video-based eye tracker at 500 Hz (EyeLink 1000, SR Research, Ontario, Canada), using the EyeLink Toolbox extensions for MATLAB (Cornelissen, Peters, & Palmer, 2002).

4.3.3 Task and Stimuli

Participants viewed a briefly presented array of three peripheral, randomly oriented Gabor patches. Their task was to remember the three stimulus orientations. During the delay interval a valid or neutral spatial retrocue was presented. At the end of the trial, participants were probed to report back one of the three orientations. Figure 4-1 provides a task schematic.

Trials began with a central fixation dot presented for 1,000-1,500 ms (randomly determined from a uniform distribution). The three target stimuli then appeared for 700 ms. After a short blank interval (70 ms) they were followed by masks (280 ms). An interval of 650 ms followed, after which a central retrocue stimulus was presented for 200 ms. After a 1,000-ms blank interval, a probe appeared at one of the three locations. In trials with spatial retrocues (valid condition), the probe appeared at the cued location. On trials with neutral retrocues (neutral condition), the probe appeared at one of the three locations randomly determined on each trial. The probe was a Gabor patch, and participants reported the target orientation by rotating the probe stimulus to the remembered target orientation by moving the computer mouse upward (counterclockwise rotation) or downward (clockwise rotation) and then clicking the mouse to confirm a response. Accuracy was emphasized, and participants were encouraged to take as

much time as they needed to complete their response. Feedback was provided for 200 ms after the response. Feedback was presented in the form of two small, white dots indicating either end of the target orientation. The subsequent ITI was 500 ms.

All stimuli appeared atop 10% contrast luminance pedestals, which were present throughout the entire experiment. The three stimuli in the memory array appeared in a triangular configuration. One stimulus was centred at 2.6° above the horizontal meridian on the vertical meridian, and the other two were located at 3.1° to the left or right from the vertical meridian and 2.6° below the horizontal meridian. Target stimuli consisted of Gabor patches (Gaussian-vignetted sinusoidal gratings) with a diameter of 3.5° of visual angle and a spatial frequency of 2 cpd of visual angle. The Gaussian envelope had a space constant of 0.44°. The range of contrast levels for the target Gabor patches varied in each of the three experiments. The stimuli used for backward masking were constructed by applying a Gaussian-vignette to the convolution of 90% contrast square-wave gratings at four orientations of 90°, 180°, 45°, and 135°. The cue stimulus (1.59° x 1.59° of visual angle) consisted of either a black arrow pointing to one of the three locations which indicated (with 100% validity) which of the three target stimuli was to be reported (valid condition). In the neutral condition, the cue stimulus consisted of the outline of a black rectangle. The probe orientation was randomly chosen on each trial.

In all cases, the specific contrast levels used in individual trials were equiprobable and randomly determined, and all contrasts were always the same within a given display. Retrocueing conditions were equiprobable and

randomized. The orientations of the target stimuli were randomly chosen on each trial.

4.3.4 Procedure

Participants completed the experiment seated in a dimly lit room. Before starting, they were given a series of training trials (minimum of 48), which were identical to the experimental task, until they were confident they could do the task. During the practice session only the highest contrast level was used. Rest breaks were provided after every block.

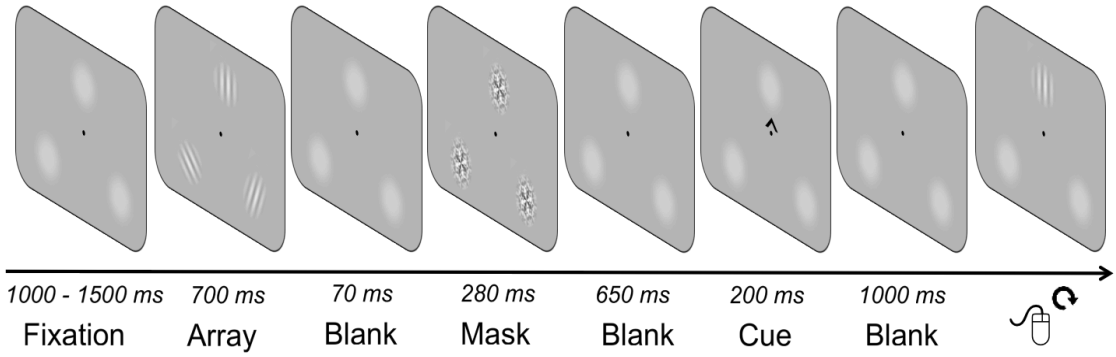

Figure 4-1. Task design for Experiments 1, 2, and 3. A central fixation point was presented for a variable duration between 1,000 – 1,500 ms (randomly determined) followed by a 700-ms target array and, after a 70-ms blank display, a 280-ms backwards-mask display. After another 650-ms blank display, a 200-ms retrocue stimulus was presented indicating which of the three stimuli they had to report at the end of the trial. After a final 1000-ms blank display, participants reported the target orientation using the computer mouse.

4.3.5 Analysis

There are considerable individual differences both in WM capacity (Adam, Mance, Fukuda, & Vogel, 2015; Vogel & Machizawa, 2004) and in contrast sensitivity (Baker, 2013). To take advantage of individual differences, I completed two sets of analysis. First, I examined the effects of attentional orienting and contrast level on recall precision at the group level

using ANOVAs. I included all participants that met the different inclusion requirements outlined below. Second, I explored how variability in WM capacity, contrast sensitivity, or both, affected performance in my task. I performed median-split analyses as well as simple linear regression analyses. In the median-split analyses, I only report effects and interactions involving the group variable, as any other effects are in principle the same as the pre-median split analysis.

4.4 Experiment 1

4.4.1 Design

The aim of the first experiment was to examine how the strength of incoming sensory information interacts with top-down orienting of spatial attention within WM to influence performance. Participants completed a total of 480 trials split into 12 blocks of 40 trials each. This yielded 60 trials per condition – valid and neutral retrocue trials at each of four possible Michelson contrast levels: 1%, 18%, 32% and 50%. The lowest contrast level was included with the aim to examine whether attentional selection in WM could raise a stimulus above threshold that otherwise would have been below-threshold for WM reports. The remaining three contrast levels were included to examine how contrast level interacted with attentional selection when all stimuli are visible (to varying degrees) and reportable.

4.4.2 Subjective Ratings

On one sixth of trials (randomly determined at the beginning of the experiment and randomly spread out across the experiment and all conditions) participants were additionally asked to rate their subjective

awareness of the stimuli using the PAS (Ramsøy & Overgaard, 2004). Ratings were performed immediately after the feedback and before the ITI. The PAS consists of four response options: 1 – stimulus not seen, 2 – weak glimpse: something was there but I do not know its orientation, 3 – almost clear image: I think I know the orientation, and 4 – clear image. Participants were instructed on how to use this scale before the beginning of the experiment. It was emphasized that ratings were to be made introspectively, relying on visual experience (Ramsøy & Overgaard, 2004), and that the ratings were to be made on the initial perceptual, and not the memory representation of the stimuli. A shortened version of the four ratings (1 – stimulus not seen, 2 – weak glimpse, 3 – almost clear image, 4 – clear image) was presented on the screen before ratings were made at the end of the trial. Responses were made by pressing the corresponding number on the keyboard.

4.4.3 Inclusion Criteria

For the individual-differences analyses, I aimed to take advantage of variability in the data set. Therefore, in the regression analysis, where I tested whether task performance would predict the magnitude of retrocueing effects, I only excluded participants if they showed chance-level performance (i.e., uniform responses, as measured with a Rayleigh test, with a threshold of $p < 0.05$) at the high contrast level (average over 18%, 32%, and 50%) in any of the retrocueing conditions. Five participants were excluded, leaving 31 participants for further analyses. Data points were removed from the linear regression analysis if they had a Cook's distance statistic greater than three times the mean Cook's distance, indicating that

these measurements played a disproportionate role in influencing the regression.

In the group-level analysis the aim was to test for differences in WM performance and retrocueing benefits with subliminal versus supraliminal stimulus presentation. Therefore, I additionally excluded participants who showed above chance-level performance at 1% contrast following neutral retrocues (n = 7). Furthermore, to reduce group variability, I only included participants who performed within 2 SD of the group mean in all conditions (i.e., not averaged over contrast). Finally, to relate the simple linear regression results to the ANOVA results, I also performed a median-split analysis using the same participant selection as for the overall ANOVA.

4.4.4 Analysis

For each trial, the reported orientation was collected. This was then measured against the orientation of the probed stimulus to derive recall error and recall precision. Recall error was defined as the angular deviation between the retrocued orientation and the reported orientation. Recall precision was computed as the trial-to-trial variability in the response error. It was calculated as the reciprocal of the SD of errors across trials. I used Fisher's definition of SD for circular data (Fisher, 1993), and subtracted the value expected by chance, so that a precision value of zero corresponds to chance-level performance.

It is important to note that the term recall precision can be used to refer to different measurements in WM studies. In some studies (e.g. Zhang & Luck, 2008) observers' WM responses are analysed assuming they reflect a mixture of distributions: a random distribution of responses where

observers guess the orientation and a von Mises (circular Gaussian) distribution centred around the target value, reflecting responses related to the target. This distribution is defined by the concentration parameter, or precision, of remembered items (k, or kappa). Some other studies add additional 'misbinding' parameters for trials in which observers mistakenly report orientations based on non-target items in the array (Bays, Catalao, & Husain, 2009). The probability of remembering a target is derived from the proportion of random responses. The concentration parameter k is used as an index of representational quality of a WM item, often also referred to as precision. Here, I use the term precision to refer to 1/SD (as defined above), while I use the term kappa to refer to concentration parameter k as an index of representational quality of items in WM. Whereas precision is a measure of overall performance, kappa measures the representational quality when an item is actually encoded, controlling for the frequency of not encoding or remembering an item. This distinction is crucial as the main hypothesis pertains to representational quality in WM.

First, I examined how retrocueing affects subliminal (1% contrast) versus supraliminal (18% - 50% contrast) performance at the group level using repeated-measures ANOVAs. I averaged over contrast levels 18% - 50% to have a measure of performance after subliminal versus supraliminal stimulus presentation. Effects of experimental parameters on recall precision were tested using repeated-measures ANOVAs with the factors of retrocue (valid, neutral) and contrast (subliminal, supraliminal). Subjective ratings were analysed using a repeated-measures ANOVA with the same factors. Results were followed up by *t* tests to guide interpretation when necessary (Bonferroni-corrected where appropriate).

Then, to examine the influence of contrast in more detail, I exploited the high variability in performance following neutral retrocues across participants. First, I performed a median-split analysis on neutral retrocue performance (recall precision, in neutral retrocue conditions averaged across 18%, 32% and 50% contrast[7]) to compare the effects of retrocueing and each of the contrast levels on recall precision in individual with high versus low WM performance. Differences between groups were tested using mixed ANOVAs with the between-subjects factor baseline performance (low, high) and the within-subjects factors retrocue (valid, neutral) and contrast (1%, 18%, 32%, 50%).

Finally, I performed linear regression to examine whether baseline performance (recall precision, in neutral retrocue conditions averaged across 18%, 32% and 50% contrast) would predict the magnitude of the retrocueing effect on recall precision.

4.4.5 Results

Data are presented in Figures 4-2, 4-3, and 4-4.

Group Analysis

Four participants had recall precision outside 2 SD from the group mean in at least one condition and were excluded from analysis to reduce group variability (total N = 21). The qualitative pattern of results remained the same whether or not these participants were excluded.

A 2-by-2 ANOVA with the factors retrocue (valid, neutral) and contrast (subliminal, supraliminal) revealed a significant retrocue-by-contrast

[7] I did not include performance at 1% in the average since the majority of individuals performed at chance at this level. The pattern of results did not change when this condition was included in the analysis.

interaction ($F(1,20) = 6.200$, $p = .022$, $\eta_p^2 = 0.237$). There was a retrocueing effect for supraliminal ($t(20) = 3.27$, $p = .004$, $d = 0.714$) but not for subliminal stimuli ($t(20) = 0.06$, $p = .950$, *ns*). Performance was at chance for both 1% conditions ($t(20) = -0.25$, $p = .806$, *ns*; $t(20) = -0.42$, $p = .680$, *ns*). Retrocue ($F(1,20) = 6.776$, $p = .017$, $\eta_p^2 = 0.253$) and contrast ($F(1,20) = 143.120$, $p < .001$, $\eta_p^2 = 0.877$) also exerted main effects reflecting that performance was better following valid retrocues and supraliminal stimulus presentation. These results suggest that retrocueing had a significant effect on recall precision only if stimuli are presented above threshold for WM report.

For subjective ratings, contrast exerted a significant main effect ($F(1,20) = 173.619$, $p < .001$, $\eta_p^2 = 0.897$). Subjective ratings were higher in the high contrast condition compared to the 1% contrast condition. None of the other effects were significant (all F's < 2.1).

To examine the effects of contrast on subjective ratings more closely, and specifically how effects of retrocueing on subjective ratings changed with contrast levels, I then split up the three high-contrast levels (18%, 32%, and 50%) and included all four levels in the analysis. This analysis revealed a main effect of contrast ($F(3,60) = 137.227$, $p < .001$, $\eta_p^2 = 2.618$) but no other effects or interactions (all F's < 2) . Follow-up *t*-tests revealed significant differences in subjective ratings between 1% contrast and 18% contrast ($t(20) = 12.550$ $p < .001$, $d = 2.739$), between 18% and 32% ($t(20) = 5.645$, $p < .001$, $d = 1.232$) but not between 32% and 50% contrast although it approached a trend ($t(20) = 1.717$, $p = 0.10$). Interestingly, this steady increase of subjective ratings with an increase in contrast was different to what I observed for recall precision when I split up the three high

contrast levels. Contrast exerted a main effect ($F(3, 60) = 102.893$, $p < .001$, $\eta_p^2 = 2.512$), but this was driven by an increase in performance from 1% to 18% ($t(20) = 11.127$, $p < .001$, d= 2.428) and a trend for an increase from 18% to 32% ($t(20) = 2.07$, $p = .051$, uncorrected, d= 0.452). Performance then remained stable and did not change with the increase to 50% contrast ($t(20) = 1.07$, ns).

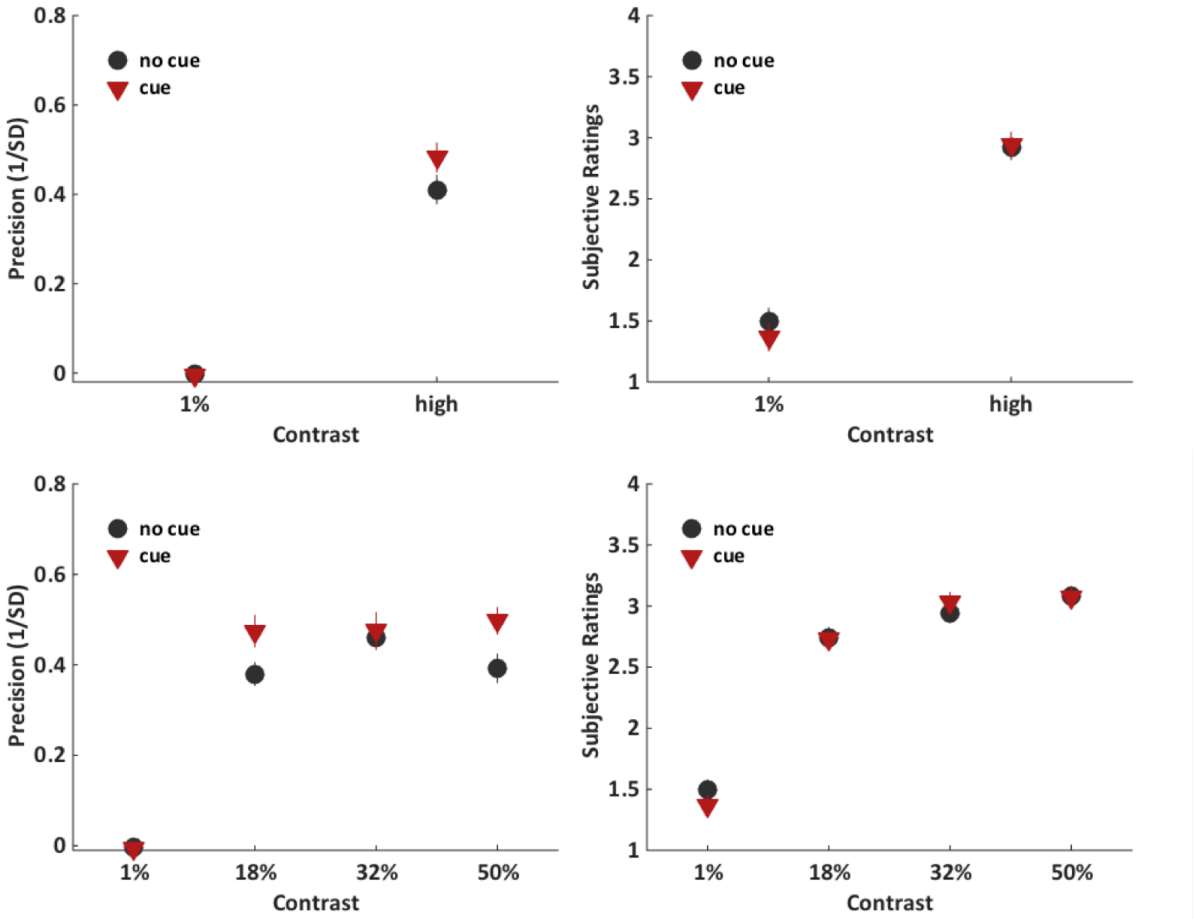

Figure 4-2. Behavioural results (precision and subjective ratings) for Experiment 1. Error bars reflect ± 1 standard error of the mean.

Median-split

To test whether baseline performance following neutral retrocues (recall precision, averaged over 18%, 32% and 50% contrast) predicted individuals' ability to benefit from the retrocues, I performed a median-split analysis on baseline performance and then compared effects of retrocueing

(valid vs neutral cue) and contrast on recall precision between groups (see Figure 4-3). This revealed a significant 3-way interaction between group, contrast, and retrocue ($F(3,57) = 4.043$, $p = .011$, $\eta_p^2 = 0.526$). To examine this effect further, I completed two-way repeated-measures ANOVAs with the levels retrocue (valid, neutral) and contrast (1%, 18%, 32%, and 50%) for each group separately. For high performers, this revealed a retrocue-by-contrast interaction ($F(3,30) = 3.189$, $p = .038$, $\eta_p^2 = 0.242$) and main effects of retrocue ($F(1,10) = 15.984$, $p = .003$, $\eta_p^2 = 0.615$) and contrast ($F(3,30) = 85.958$, $p < .001$, $\eta_p^2 = 0.89$). Recall precision improved, or had a tendency to improve, following valid retrocues but only for supraliminal contrast levels (18%: $t(10) = 4.740$, $p = .001$, $d = 1.499$; 32%: $t(10) = 1.822$, $p = .098$, $d = 0.549$; 50%: $t(10) = 1.88$, $p = .089$, $d = 0.567$) and not at 1% contrast ($t(10) = 0.248$, *ns)*.

For low performers there was a significant retrocue-by-contrast interaction ($F(3,27) = 4.203$, $p = .021$, $\eta_p^2 = 0.955$). Recall precision improved following valid retrocues but only at 50% contrast ($t(9) = 3.146$, $p = .0118$, $d = 0.995$; all other t's < 1.5). In other words, for high performers retrocueing effects were consistent when stimuli were presented above threshold, but for low performers retrocueing effects only emerged at the highest contrast level.

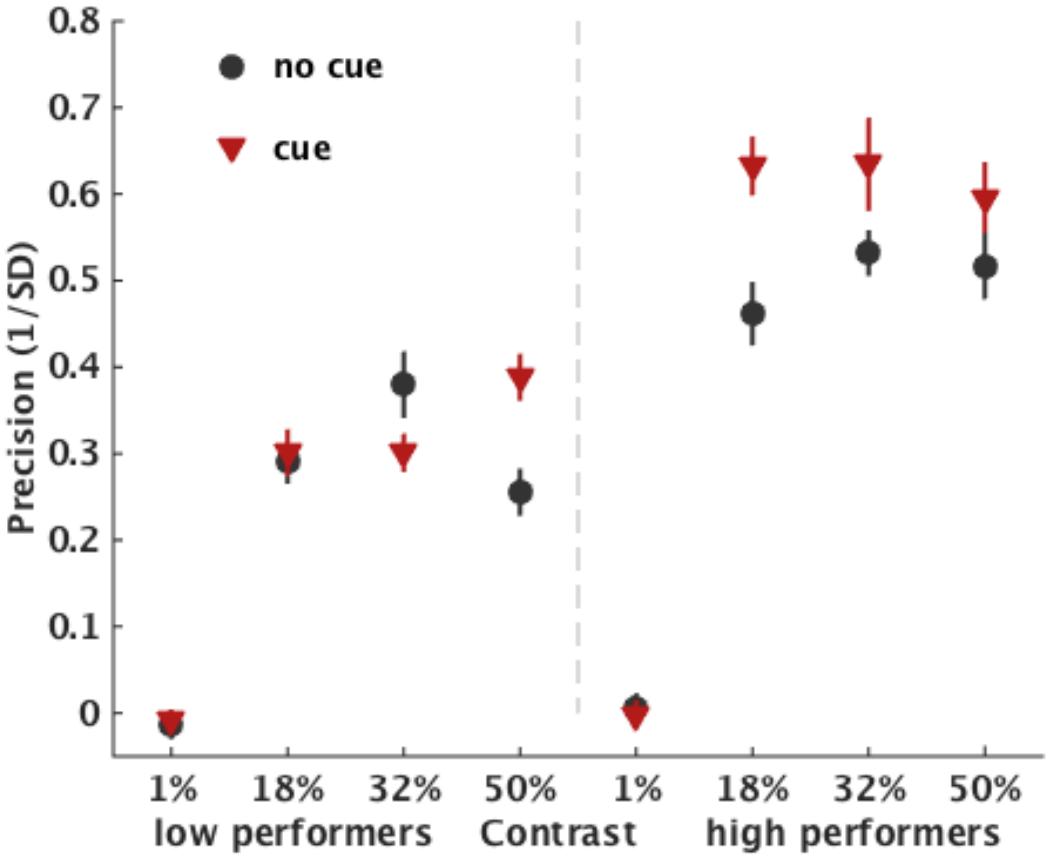

Figure 4-3. Recall precision (1/SD) in Experiment 1 for each condition for high-performance (right) and low-performance (left) groups separately. Error bars reflect ± 1 standard error of the mean.

Linear regression

Linear regression analyses showed that baseline performance (recall precision averaged across 18%, 32% and 50% contrast in neutral retrocue condition) predicted the magnitude of the retrocueing effect on recall precision with an R^2 of .281 (adjusted R^2 = .253), $F(1,25) = 9.792$, $p <$ .005, β = .531). Specifically, the better observers' baseline performance was, the larger the retrocueing effect was on recall precision.

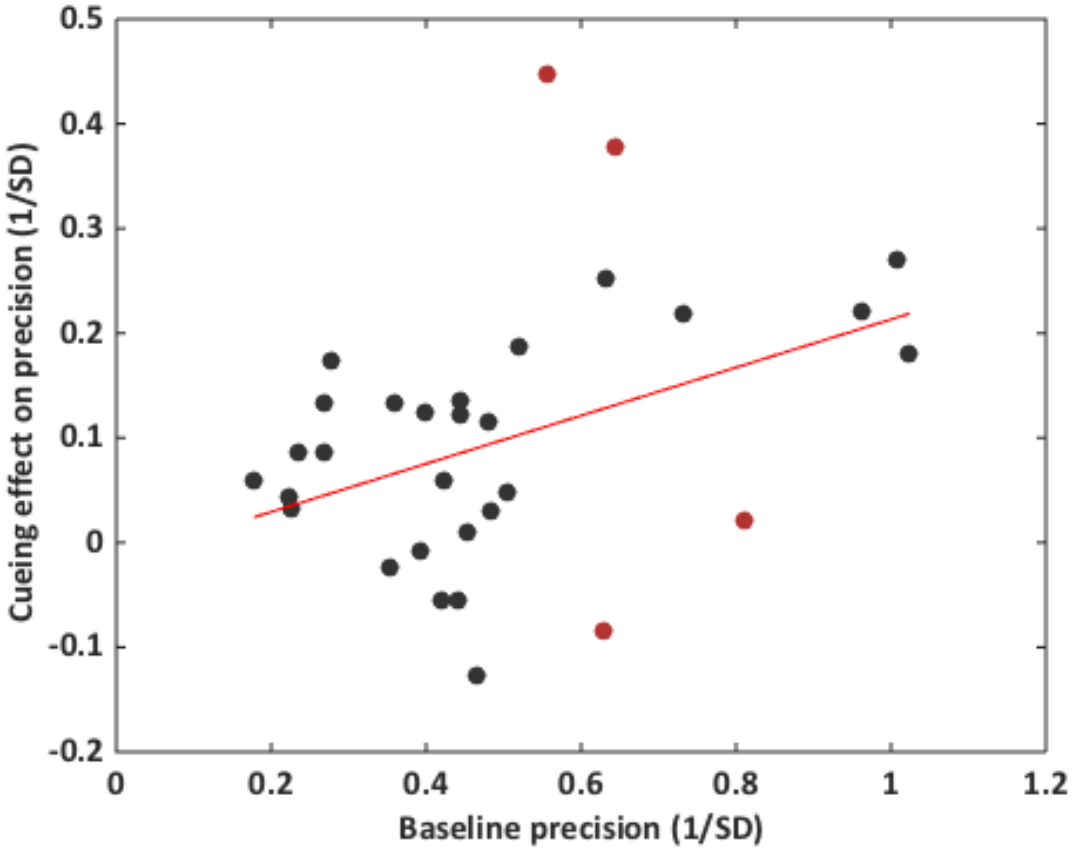

Figure 4-4. Baseline performance (precision, averaged over contrast levels 18%, 32% and 50% in the neutral retrocue condition) predicts magnitude of the retrocueing effect on recall precision, $p < .005$. Outliers (n = 4, shown in red) were removed when they had a Cook's distance statistic greater than three times the mean Cook's distance. The exclusion of outliers lead to a quantitative change in the value of beta coefficients but it did not qualitatively change the results.

4.4.6 Interim Discussion

The results of Experiment 1 show that top-down attentional orienting during WM maintenance has an effect on recall precision, adding to the growing literature demonstrating the effect of 'retrocues' on WM task performance (Griffin & Nobre, 2003; Landman et al., 2003; Lepsien & Nobre, 2006; Lepsien, Griffin, Devlin, & Nobre, 2005; Makovski, 2012; Rerko et al., 2014; Rerko & Oberauer, 2013; Sligte et al., 2008; Williams et al., 2013). The results also reveal that retrocueing effects in WM were

modulated by contrast: significant retrocueing effects only emerged for supraliminal stimuli presented at high contrast levels, but not for subliminal stimuli presented at 1% contrast.

I further examined the effects of different contrast levels in relation to the high variability in WM performance following neutral retrocues. A median-split analysis revealed that high performers showed cueing effects, or a tendency for retrocueing effects, at all-contrast levels that allowed for above-chance performance, while low performers only showed retrocueing benefits at the highest contrast level. These findings show that the quality of incoming sensory information interacts with individual WM capacity to shape WM performance. Furthermore, the findings suggest that retrocueing effects emerge only once WM representations are of a minimal quality.

In the present study it was not possible to determine which components of WM recall were improved following valid retrocues, because I looked at the overall error distribution in the form of recall precision, and did not separate WM responses into different components. Specifically, my results could be explained by changes in the quality of a WM representation, its likelihood of being remembered, or mistakenly reporting a distracter item instead of the target item. The number of trials per condition prevented me from doing reliable modelling to estimate each of these components and I cannot know how retrocueing affected different aspects of WM recall at different levels of contrast. This will be explored in Experiments 2 and 3.

I also found that sub-threshold stimuli did not benefit from attentional orienting: recall performance remained at chance following valid retrocues. Attention has long been considered to play an integral role in conscious

experience (Baars, 1988; Posner, 1994). According to Baars' Global Workspace Theory (1988), neural activation elicited by a stimulus needs to be amplified top-down through attentional mechanisms for a stimulus to be consciously experienced. According to this account, attention has a direct role in shaping and strengthening unconscious representations. Indeed, (Carrasco et al., 2004) have shown that when the contrast of a stimulus is just below the threshold of visibility, attending to the stimulus can increase its salience rendering it visible when otherwise it would have remained invisible. My findings, using retrocues, could reflect one of two things. One possibility is that mechanisms are different for WM. For example, weak, subliminal representations may not remain long enough to be boosted by attention. Varying the time point at which the cue is presented during the WM delay period could test this hypothesis. Alternatively, the neural activity elicited by the 1% contrast stimulus was too weak to exceed a threshold even after attentional boosting, for example if the stimulus was not perceived or encoded at all.

Finally, I also observed a dissociation between subjective ratings of awareness and objective performance: while subjective ratings steadily increased with the increase in contrast, recall precision did not. Bona, Cattaneo, Vecchi, and Soto (2013) also reported a dissociation between objective measures of WM performance and subjective ratings of the vividness of the memory item. Specifically, objective performance was impaired by markedly different distracters presented at all levels of visibility, while subjective ratings were impaired by all distracter types but only when they were presented subliminally. Here, I asked participants to rate the visibility of the perceptual stimuli and not their memory of the stimuli;

however, since I did not implement any control measures, I cannot know whether this is indeed what participants did. Possibly, participants were unable to differentiate the visibility of initial stimulus presentation from their memory of the stimulus' visibility. Alternatively, subjective ratings for perceptual stimuli and memory of perceptual stimuli are both dissociable from objective WM performance. I also did not find an influence of retrocueing on subjective ratings, however, given the small trial numbers and the fact that I only included subjective ratings as a control of conscious awareness to complement objective performance, I will not interpret this any further here.

4.5 Experiment 2

4.5.1 Design

The aim of this follow-on experiment was to separate the effects of retrocueing and contrast on different distributions making up the overall WM error distribution. To this end, I doubled the number of trials so I could perform mixture modelling to estimate the different sources of error (Bays et al., 2009; see below for details). Participants completed a total of 480 trials split into 12 blocks of 40 trials each (120 trials per condition). Only two contrast levels were used for the stimulus array - 18% and 50% - resulting in more trials per condition for modelling of the behavioural results.

4.5.2 Analysis

For each trial, the reported orientation was collected. This was then measured against the orientation of the probed stimulus to derive recall error and recall precision. I modelled different sources of errors contributing

to memory performance by applying a probabilistic model to the data, as already described in Chapter 3 (Bays et al., 2009). The model attributes the overall distribution of responses to a mixture of three possible sources of errors on each trial: 1) a von Mises distribution (circular analogue of the Gaussian) centred on the target direction, 2) a uniform distribution of error corresponding to random responses, and 3) von Mises distributions centred around each of the non-target orientations in a given stimulus display, i.e. a proportion of errors that corresponds to participants mistakenly reporting an orientation based on a non-target item. This model is described by the following equation:

$$p(\hat{\theta}) = a\Phi_K(\hat{\theta} - \theta) + \beta \frac{1}{m} \sum_i^m \Phi_k(\hat{\theta} - \varphi_i) + \gamma \frac{1}{2\pi}$$

where θ is the actual orientation of the target, $\hat{\theta}$ is the reported orientation, Φ_K is the von Mises distribution with mean zero and concentration parameter k. The concentration parameter k corresponds to the variability of target recall, where greater k corresponds to lower variability in the distribution. I refer to it as kappa and use it as an index of representational quality of WM items. The probability of reporting the targets, or the probability of remembering the target item, is given by a. The probability of reporting a distracter is given by β which I use as an index of misbinding, andφ_1, φ_2, …, φ_m are the orientations of the m non-target items. The probability of randomly responding (or guessing) is given by $\gamma = 1 - a - \beta$.

Maximum-likelihood estimates of each parameter (a, β, γ, and k) were obtained by using an expectation-maximization algorithm separately for each participant and experimental condition (Bays et al., 2009).

As in Experiment 1, I explored individual differences in performance by performing a median-split and a linear regression analysis. All procedures were identical to before, except that performance was averaged over contrast levels 18% and 50% in neutral retrocue conditions for an estimate of baseline performance. Additionally, I also performed median-split analyses on the mixture-model parameters. Specifically, I performed median-split analysis on baseline kappa (averaged across 18% and 50% in the neutral retrocue condition) to compare the effect of retrocueing and contrast on kappa, misbinding, target – response probability and guess rates in individuals with high versus low WM performance. I split the data based on kappa for all parameters because I were interested specifically in individual differences in the representational quality WM items, and how differences in representational quality impacts on effects of retrocueing and contrast. I included all participants that met the inclusion requirements outlined above.

Second, I examined the effects of experimental parameters on recall precision, kappa, misbinding, the probability of target recall, and guess rate at the group level using ANOVAs with the factors of retrocue (valid, neutral) and contrast (18%, 50%). Additional *t* tests (Bonferroni corrected where appropriate) were run when necessary.

4.5.3 Results

Behavioural data are presented in Figures 4-5, 4-6, 4-7, 4-8 and 4-9.

Group Analysis

After excluding participants with performance outside 2 SD from the group mean, the following numbers of participants were entered into the

analyses: recall precision N = 27; target recall probability N = 28; misbinding N = 25; guess rate = 27 and kappa N = 24. The qualitative pattern of results remained the same whether or not these participants were excluded.

Retrocue validity exerted a significant main effect on recall precision ($F(1,26) = 23.04$, $p < .001$, $\eta_p^2 = 0.470$). There was no significant effect of contrast ($F(1,26) = 2.72$, $p = .111$), and the two factors did not interact ($F(1,26) = 1.449$, $p = .239$). These results replicated the effect of retrocue validity reported in Experiment 1, where recall precision was better following valid compared to neutral retrocues, and this did not change with stimulus contrast as long as stimuli were presented above threshold.

For kappa, there was a significant interaction between retrocue and contrast ($F(1,23) = 14.647$, $p = .001$, $\eta_p^2 = 0.389$). The main effects of retrocue and contrast were also significant ($F(1,23) = 7.721$, $p = .011$, $\eta_p^2 = 0.251$ and $F(1,23) = 14.111$, $p = .001$, $\eta_p^2 = 0.380$, respectively). Kappa was higher following valid compared to neutral retrocues but only at 50% contrast ($t(23) = 3.99$, $p < .001$, $d = 0.82$) and not at 18% contrast ($t(23) = 0.37$ $p = .718$, $d = 0.07$).

For target recall probability there was a significant main effect of retrocue ($F(1,27) = 8.87$, $p = .006$, $\eta_p^2 = 0.247$). Observers reported the target orientation significantly more often following valid compared to neutral retrocues. The main effect of contrast was not significant ($F(1,27) = 0.004$, $p = .952$) and the two factors did not interact ($F(1,27) = 0.138$, $p = .713$).

There were no significant main effects or interactions on misbinding (all F's < 2.1), but there was a strong trend for a main effect of retrocue on guess rate ($F(1,26) = 3.77$, $p = .063$, $\eta_p^2 = 0.1266$). The guess rate tended

to be lower following valid retrocues relative to neutral retrocues. Neither the main effect of contrast nor the retrocue by contrast interaction was significant (all F's < 1).

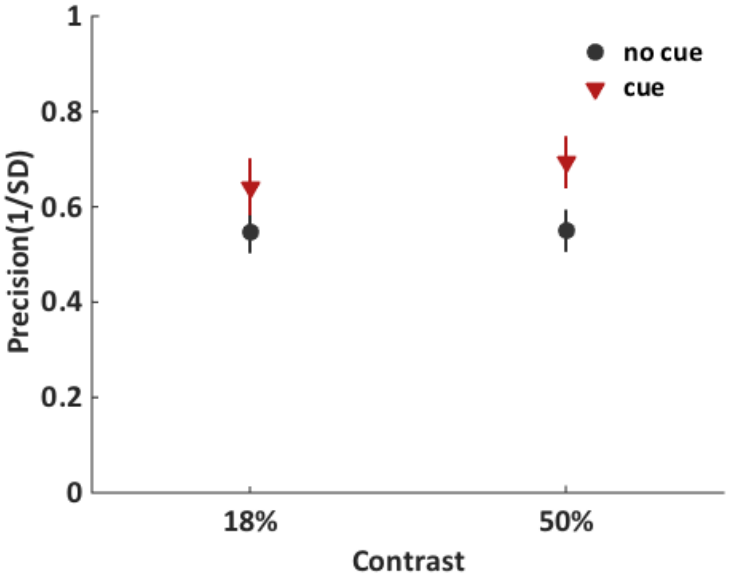

Figure 4-5. Recall precision for each condition separately of Experiment 2. Error bars reflect ± 1 standard error of the mean.

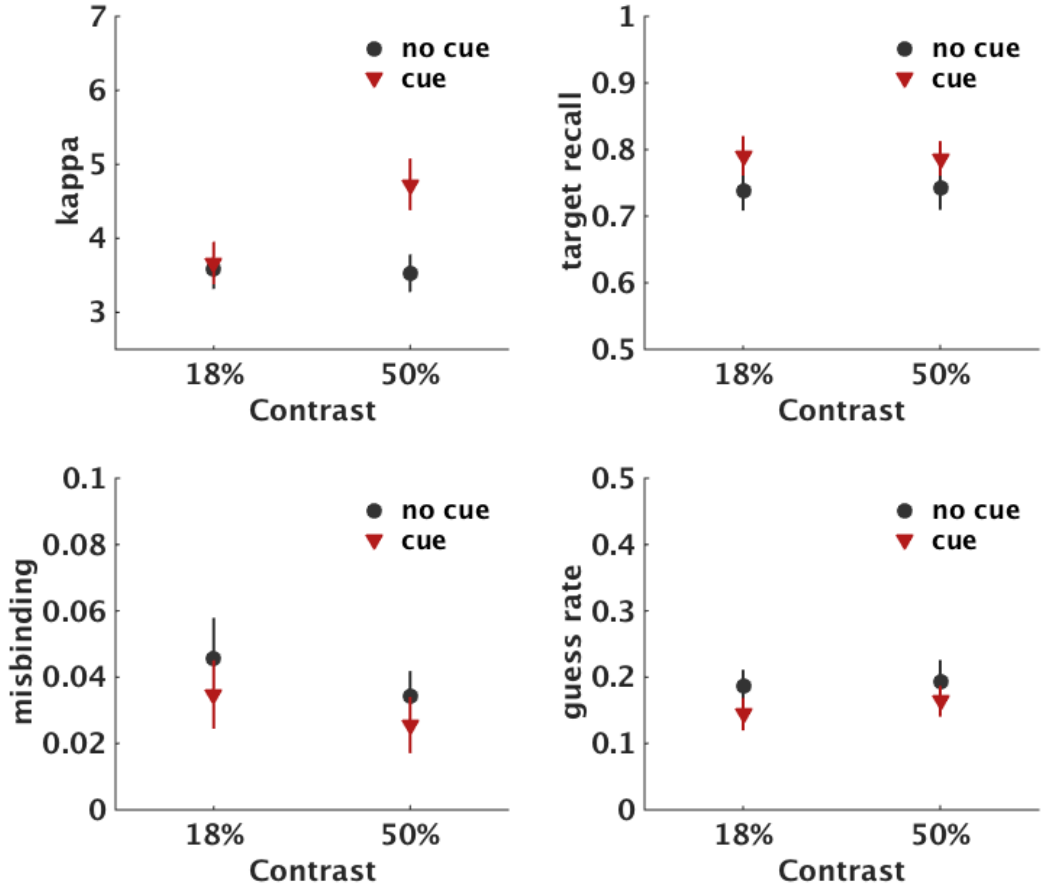

Figure 4-6. Modelling parameters for each condition separately of Experiment 2. Error bars reflect ± 1 standard error of the mean.

Median-split

I performed the same median-split analysis as in Experiment 1. This revealed a significant interaction between group and retrocue ($F(1,25) = 9.467$, $p = .005$, $\eta_p^2 = 0.275$), reflecting that retrocueing effects (valid minus neutral retrocue) tended to be larger in the high group compared to the low group (t-test comparing retrocueing effects averaged over contrast levels between groups, $t(25) = 1.824$, $p = .08$, $d = 0.358$; M_{high}: 0.162 ± 0.042, M_{low}: 0.075 ± 0.02).

To test whether baseline kappa (averaged over 18% and 50% contrast in neutral retrocue condition) predicted retrocueing effects on

kappa, target-response probability, misbinding and guess rate, I performed a median-split analysis on baseline kappa and then compared effects of retrocueing and contrast on kappa, target-response probability, misbinding, and guess rate between groups.

For kappa, there was a marginally significant three-way interaction between group, retrocue, and contrast ($F(1,22) = 4.63$, $p = .043$, $\eta_p^2 = 0.174$). Follow-up t-tests revealed that there was a significant retrocueing effect only at 50% contrast and for high performers ($t(11) = 4.376$, $p = .001$, uncorrected, $d = 1.263$).

For probability of target recall, guess rate, and misbinding there were no effects or interactions involving group (all F's < 2), although for misbinding the main effect of group approached a trend toward overall lower rates of misbinding for high performers ($F(1,23) = 2.892$, $p = .102$).

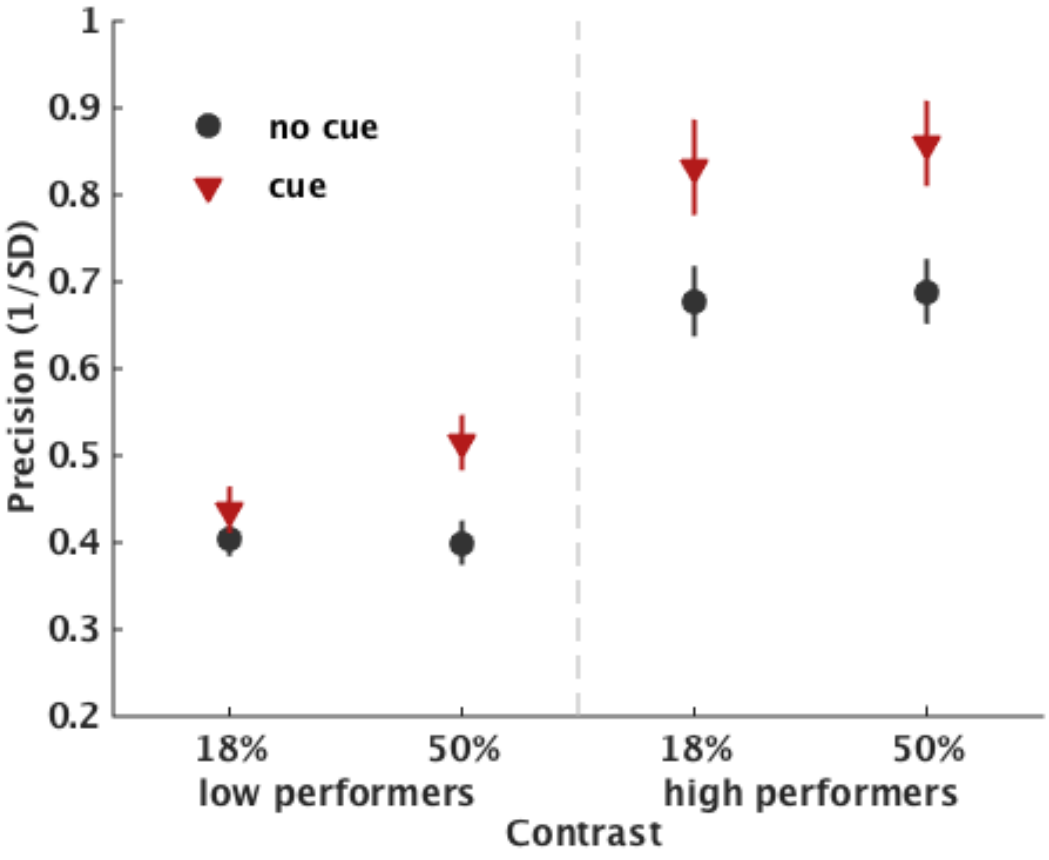

Figure 4-7. Recall precision in Experiment 2 for each condition for high-performance (right) and low-performance (left) groups separately. Error bars reflect ± 1 standard error of the mean.

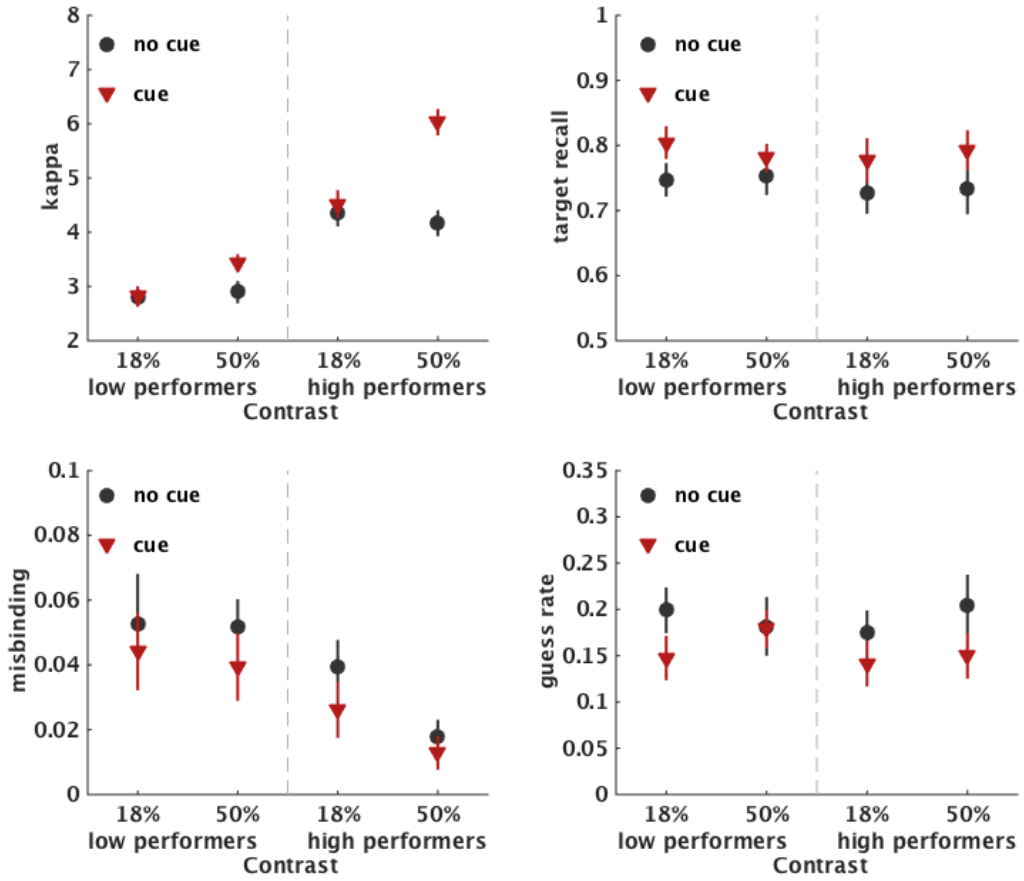

Figure 4-8. Modelling parameters in Experiment 2 for each condition for high-performance (right) and low-performance (left) groups separately. Error bars reflect ± 1 standard error of the mean.

Linear regression

Linear regression analyses showed that baseline performance (recall precision, averaged across 18% and 50% contrast in neutral retrocue condition) predicted the magnitude of the retrocueing effect on recall precision with an R^2 of .254 (adjusted R^2 = .224), $F(1,25) = 8.490$, $p = .007$, $\beta = .503$. Individuals with better baseline precision benefitted more from valid retrocueing.

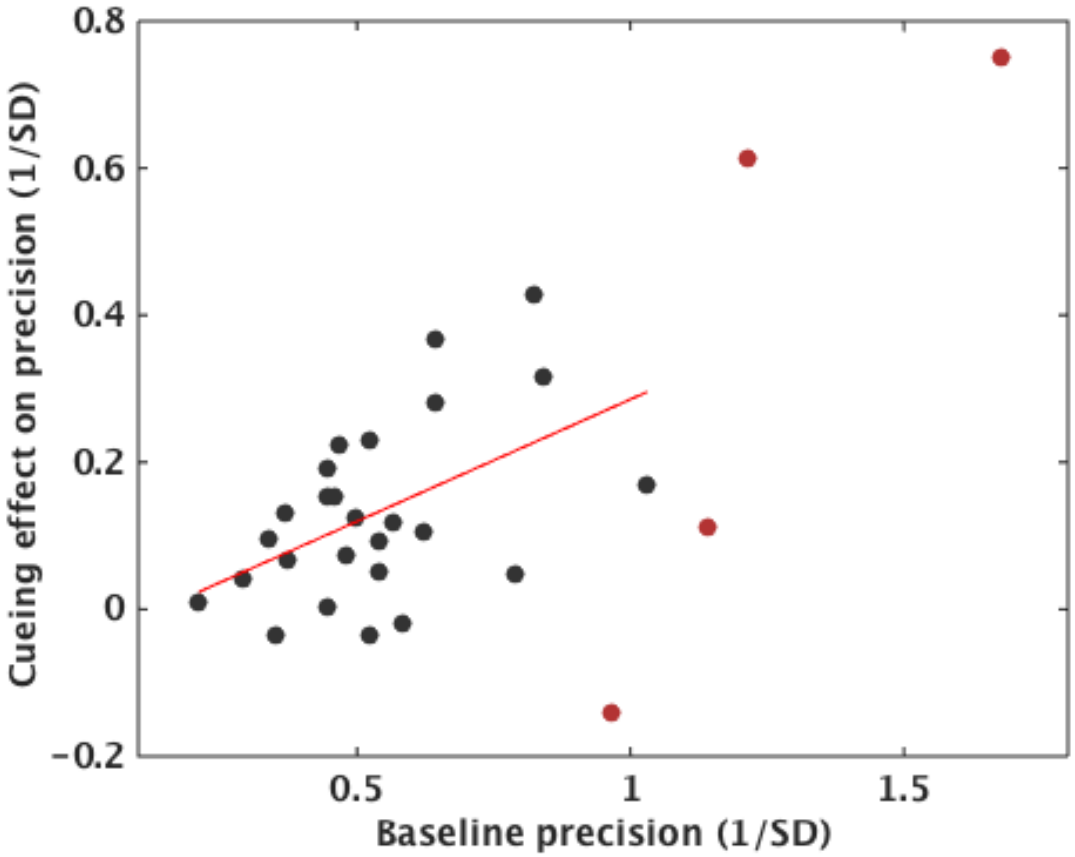

Figure 4-9. Baseline performance (precision, averaged over contrast levels 18%, and 50% in the neutral retrocue condition) predicts magnitude of the retrocueing effect on recall precision, $p < .01$. Outliers (n = 4, shown in red) were removed when they had a Cook's distance statistic greater than three times the mean Cook's distance. The exclusion of outliers lead to a quantitative change in the value of beta coefficients but it did not qualitatively change the results.

4.5.4 Interim Discussion

As in Experiment 1, the results showed that recall precision was improved following valid retrocues, and this effect was significantly modulated by individual differences in WM performance, i.e. by variability in recall precision following neutral retrocues. High-performing individuals showed a stronger retrocueing effect than low-performing individuals, and baseline recall precision predicted the magnitude of the retrocueing effects. The retrocueing benefit was not modulated by contrast levels. This supports

findings from Experiment 1, where I also found that baseline precision is not strongly modulated by stimulus contrast once Gabor patches were presented at a contrast level that enabled observers to perform above chance-level.

The ability to model the data with the increased number of trials allowed me to understand what variables contributed to the changes in performance. Specifically, I showed that the overall improvement in recall precision following valid retrocues was driven by an increase in the representational quality of WM items (kappa) at high contrast levels, and increase in remembering the target item, and a decrease in guessing, and. The proportion of misbinding errors was not modulated by retrocue validity or contrast.

Furthermore, I found that individual differences in baseline performance not only affected retrocueing effects quantitatively, but also qualitatively: the performance benefit following valid compared to neutral retrocues was more pronounced for individuals with high baseline recall precision, and an increase in kappa following valid retrocues was only observed for these high-performing individuals. Low-performers only benefitted from retrocues through an increase in the probability of remembering an item, and a tendency for a reduced guess rate. Individual differences in baseline performance did also somewhat modulate overall misbinding rates.

In Experiments 1 and 2 I did not observe an influence of stimulus contrast (at supraliminal levels) on baseline performance. It seems that the contrast levels used in both experiments were beyond any effects of facilitation, and performance never approached the lower asymptote of the

psychometric curve. Instead, performance was at chance at 1% contrast and had already reached a plateau at 18% contrast. One possible interpretation is that memory for perceived orientations is categorical, and a stimulus is remembered at a representational quality independent of its quality during sensory encoding as long as it is presented above threshold. Alternatively, it is still possible that memory quality is affected by the sensory quality of the stimulus during encoding, but that the contrast levels used in the previous two studies might have missed the range of any facilitation effect. Similarly, relative patterns of misbinding and kappa, and retrocueing effects on these parameters, may change as performance approaches the lower asymptote considering that both kappa and misbinding are modulated by baseline performance.

Experiments 1 and 2 gave me no leverage on the question of whether graded effects of contrast might be observed with low contrast levels. To explore whether retrocues rely on different mechanism as stimulus contrast transitions from subliminal to at-threshold and on to clearly visible, and to investigate how such changes might affect memory recall following neutral retrocues, I examined in more detail performance with low contrast levels (between 1% and 18%) in greater detail in Experiment 3. An alternative to using set contrast levels would have been to titrate contrast levels based on individual psychophysical thresholds. I chose not to do this because it is not straightforward to translate results from a standard forced-choice staircasing procedure to the WM task used here[8]. I therefore decided

[8] Pilot data collected as part of this thesis revealed that there were differences in the absolute contrast levels required for above-chance performance in a standard forced-choice staircasing procedure and above-chance performance in the WM task.

to capitalise on the natural variability in performance in neutral-retrocue trials.

4.1. Experiment 3

4.5.5 Design

Building on the findings of the first two experiments, the third experiment samples the contrast-level space in between the subliminal and low-contrast levels more extensively (3%, 10%, and 18%). A second aim of the Experiment was to characterise patterns of kappa and misbinding at lower levels of contrast. As in Experiment 2, large trial numbers were used to enable modelling of the data to reveal the pattern of effects across kappa, target response probability, misbinding and proportion of guessing. Participants completed a total of 600 trials split into 15 blocks of 40 trials each (100 trials per condition). Eye gaze was tracked to ensure that differences between conditions were not related to differences in eye-movements.

4.5.6 Analysis

The analysis approach followed that of Experiment 2. I performed analyses on individual differences in performance by median-split and linear regression analyses. Performance was averaged over contrast levels 3%, 10% and 18% in neutral retrocue conditions for an estimate of baseline performance.

I examined effects of experimental parameters on recall precision, kappa, misbinding, the probability of target recall and guess rate using repeated-measures ANOVAs with factors of retrocue (valid, neutral) and

contrast (3%, 10%, 18%) with follow-up *t* tests where needed (Bonferroni corrected where appropriate). When eye-tracking data were available, trials in which eyes moved more than 2 degrees from fixation were removed from all analyses (~6% trials). For seven participants eye-tracking data were not available due to technical difficulties during recording.

4.5.7 Results

Behavioural data are presented in Figures 4-10, 4-11, 4-12, 4-13 and 4-14.

Group Analysis

After excluding participants with performance outside 2 SD from the group mean the following numbers of participants were entered into the analyses: recall precision N = 31; target recall probability N = 30; misbinding N = 28; guess rate N = 29 and kappa N = 28. The qualitative pattern of results remained the same whether or not these participants were excluded.

For recall precision, there was a significant retrocue-by-contrast interaction ($F(2,60) = 3.878$, $p = .0028$, $\eta_p^2 = 0.114$). To examine the interaction more closely, I ran follow-up *t* tests comparing recall precision following valid-retrocue and neutral-retrocue trials at each of the three contrast levels. These revealed significant retrocueing effects at contrast levels 3% ($t(30) = 3.043$, $p = .005$, $d = 0.547$), at 10% ($t(30) = 4.614$, $p < .0001$, $d = 0.829$) and at 18% ($t(30) = 3.944$, $p < .001$, $d = 0.708$). However, retrocueing effects were significantly smaller at 3% contrast compared to 10% contrast ($t(30) = 2.780$, $p = .009$, $d = 0.499$) and marginally different to retrocueing effects at 18% ($t(30) = 1.906$, $p = .066$, d

= 0.342). Retrocueing effects at 10% were not different to retrocueing effects at 18% (t(30) = 0.926, p = .362, ns; $CueingEffect_{3\%}$ = 0.0633 ± 0.01; $CueingEffect_{10\%}$ = 0.16 ± 0.02; $CueingEffect_{18\%}$ = 0.126 ± 0.02). In other words, retrocueing effects increased from 3% to 10% and then remained stable. I also observed significant main effects of retrocues (F(1,30) = 29.447, p < .001, η_p^2 = 0.495) and contrast (F(2,60) = 6.506, p = .005, η_p^2 = 0.178).

For kappa none of the effects were significant (all F's < 1.2).

For target recall probability there was a significant main effect of retrocue (F(1,29) = 13.994, p = .001, η_p^2 = 0.326). Observers reported the target orientation significantly more often following valid compared to neutral retrocues. The main effect of contrast and the contrast-by-retrocue interaction were not significant (F(2,58) = 0.683, ns, and (F(2,58) = 1.138, ns; respectively).

For misbinding there was a significant main effect of retrocue (F(1,27) = 10.386, p = .003, η_p^2 = 0.278). Observers reported the non-target orientation significantly more often following neutral compared to valid retrocues. The main effect of contrast (F(2,54) = 0.520, ns) and the retrocue-by-contrast interaction (F(2,54) = 0.607, ns) were not significant.

For guess rate, the main effect of retrocueing was marginally significant (F(1,28) = 4.069, p = .053, η_p^2 = 0.127) reflecting that guess rate decreased following valid retrocueing. The main effect of contrast (F(2,56) = 0.685, ns) and the retrocue-by-contrast interaction (F(2,54) = 0.191, ns) were not significant.

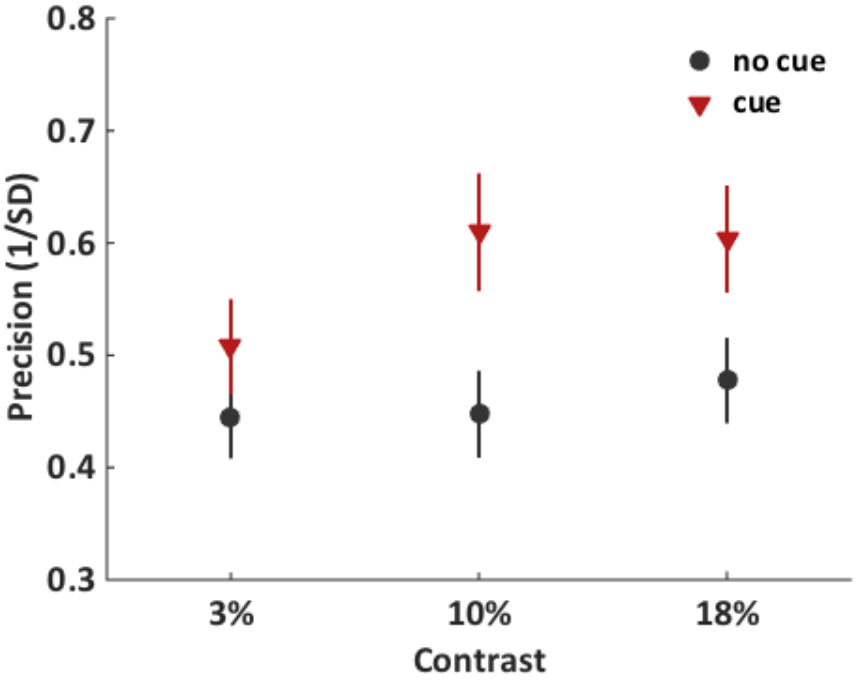

Figure 4-10. Recall precision for each condition separately of Experiment 3. Error bars reflect ± 1 standard error of the mean.

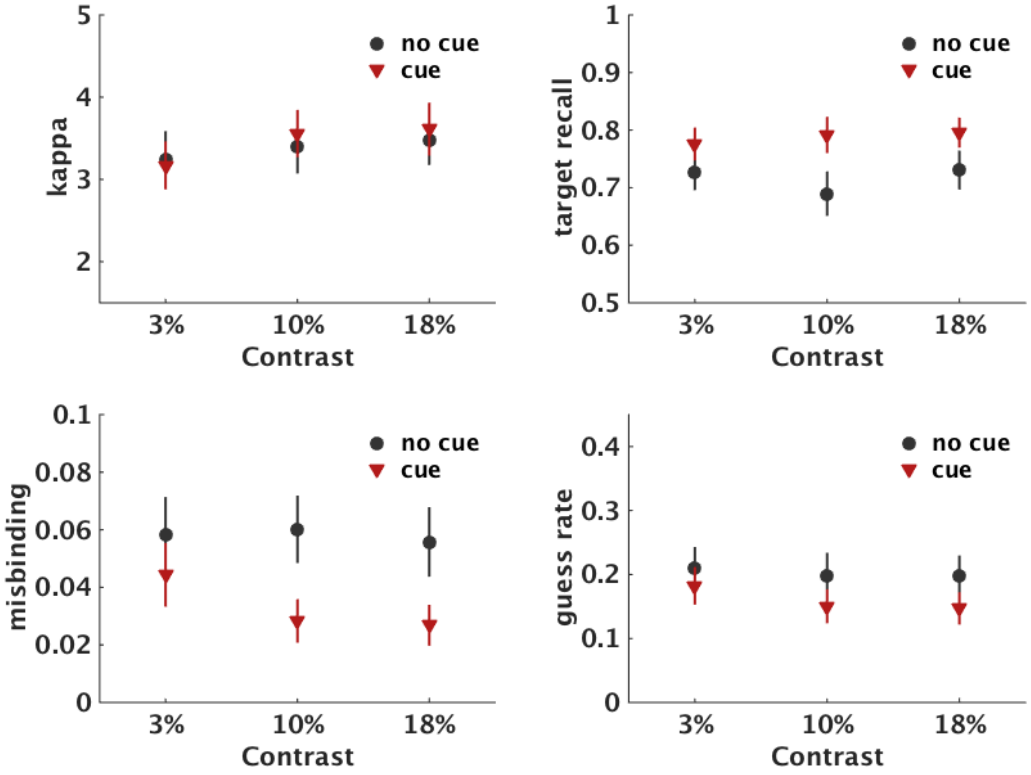

Figure 4-11. Modelling parameters for each condition separately of Experiment 3. Error bars reflect ± 1 standard error of the mean.

Median-split

To test whether baseline performance (recall precision, averaged over 3%, 10% and 18% contrast in neutral retrocue condition) predicted individuals' ability to benefit from the retrocues, I performed a median-split analysis on baseline performance and then compared effects of retrocueing and contrast on recall precision between groups. This revealed a strong trend for an interaction between group and retrocue ($F(1,28) = 4.121$, $p = .052$, $\eta_p^2 = 0.128$). Valid retrocueing benefits tended to be larger in the high-performance group compared to the low-performance group ($t(28) = 2.030$, $p = .052$, $d = 0.377$; M_{high}: 0.153 ± 0.036, M_{low}: 0.070 ± 0.020). There was also a significant group-by-contrast interaction ($F(2,56) = 3.793$, $p = .0285$, $\eta_p^2 = 0.119$). Contrast levels affected low performers ($F(2,28) = 3.633$, $p = .040$ $\eta_p^2 = 0.206$) and there was a strong trend for an effect of contrast levels for high performers ($F(2,28) = 3.794$, $p = .053$, $\eta_p^2 = 0.213$). There were also significant main effects of retrocue ($F(1,28) = 29.68$, $p < .001$, $\eta_p^2 = 0.515$), group ($F(1,28) = 101.5$, $p < .001$, $\eta_p^2 = 0.784$), and contrast ($F(2,56) = 6.598$ $p = .003$, $\eta_p^2 = .191$), reflecting that recall precision tended to be higher in the high-contrast condition, for high-performers, and following valid retrocues.

To test whether baseline kappa (averaged over 3%, 10%, and 18% contrast in neutral retrocue condition) predicted retrocueing effects on kappa, target-response probability, misbinding and guess rate; I performed a median-split analysis on baseline kappa and then compared effects of retrocueing and contrast on kappa, target-response probability, misbinding, and guess rate between groups.

For kappa, group ($F(1,26) = 35.2$, $p < .001$, $\eta_p^2 = 0.575$) exerted a main effect reflecting that kappa was higher in the high group. None of the other effects or interactions were significant (all F's < 2).

For probability of target recall, retrocue ($F(1,28) = 13.855$, $p < .001$, $\eta_p^2 = 0.331$) exerted a main effect reflecting that target-recall probability was higher following valid retrocues. None of the other effects or interactions were significant (all F's < 2).

For guess rate, there was a trend for a main effect of retrocue ($F(1,27) = 3.310$, $p = .08$, $\eta_p^2 = 0.109$) reflecting that guess rate was lower following valid retrocues. None of the other effects or interactions were significant (all F's < 2).

For misbinding, there was a significant three-way interaction between group, retrocue and contrast ($F(2,52) = 4.055$, $p < .05$, $\eta_p^2 = 0.135$). Follow-up t-tests revealed that there was a significant cueing effect only at 18% contrast for low performers ($t(13) = 2.561$, $p = .0237$, uncorrected, $d = 0.685$), and at 10% contrast for high performers ($t(13) = 2.567$, $p = .0234$, uncorrected, $d = 0.686$). In addition, retrocue ($F(1,26) = 10.21$, $p = .004$, $\eta_p^2 = 0.282$) exerted a main effect reflecting that misbinding was lower following valid retrocues. There was a also a significant group by contrast interaction ($F(2,52) = 4.243$, $p = .020$, $\eta_p^2 = 0.140$), reflecting that groups differed in misbinding rates at 10% contrast ($F(1,13) = 4.725$, $p = .049$, $\eta_p^2 = 0.267$) but not at 3% or 18% contrast (all F's < 2.2).

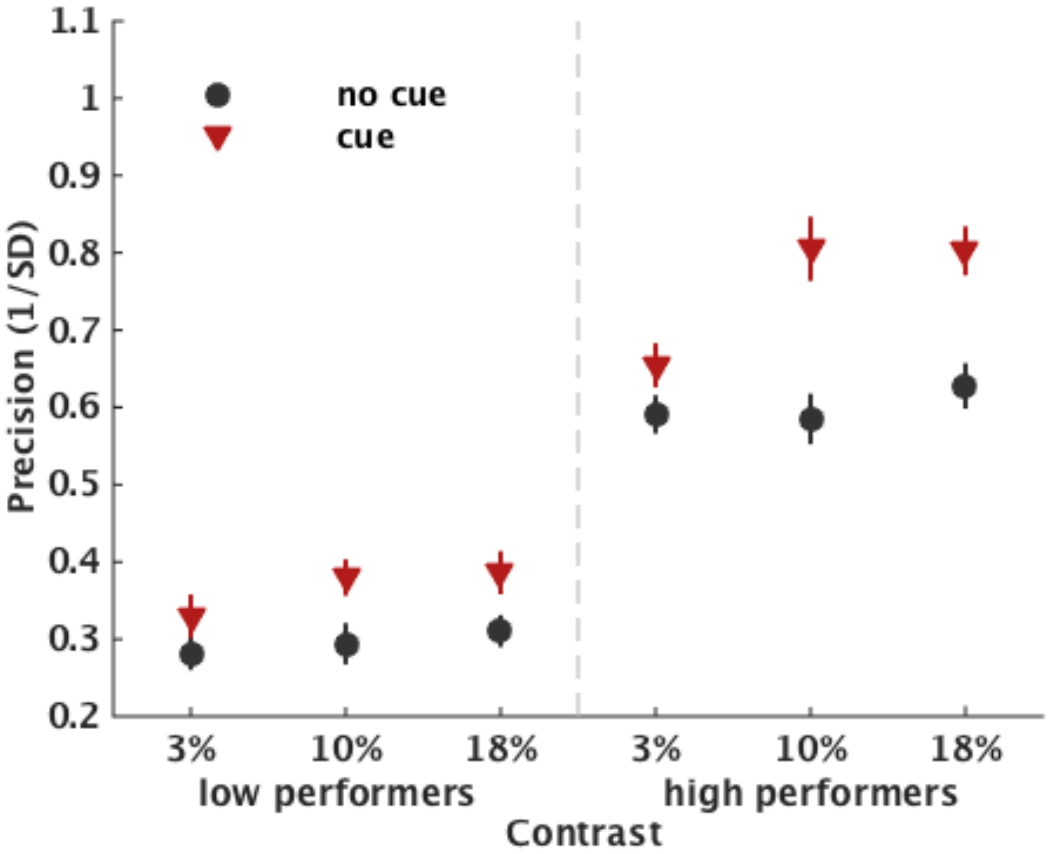

Figure 4-12. Recall precision in Experiment 3 for each condition for high-performance (right) and low-performance (left) groups separately. Error bars reflect ± 1 standard error of the mean.

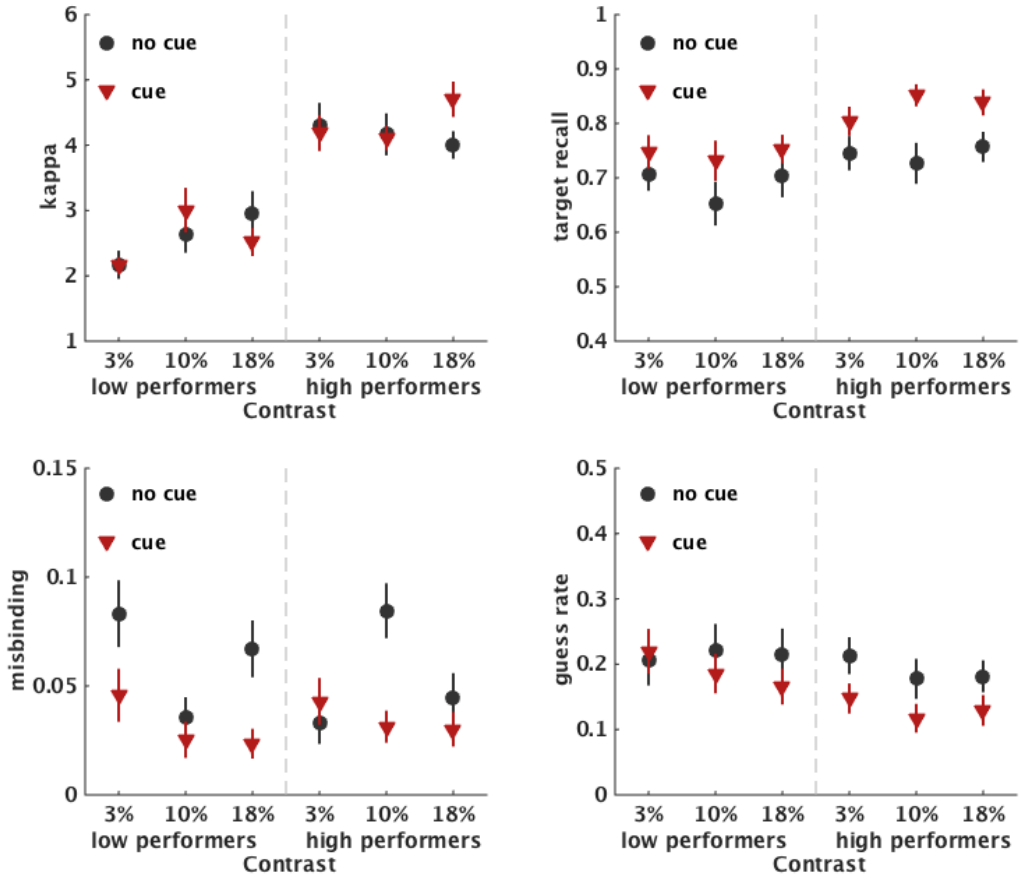

Figure 4-13. Modelling parameters in Experiment 3 for each condition for high-performance (right) and low-performance (left) groups separately. Error bars reflect ± 1 standard error of the mean.

Linear regression

Linear regression analyses showed a trend for baseline performance (recall precision, averaged across 3%, 10%, and 18% contrast in neutral retrocue condition) predicting the magnitude of the retrocueing effect on recall precision an R^2 of .087 (adjusted R^2 = .055), $F(1,29) = 2.755$, p = .108, $\beta = .295$.

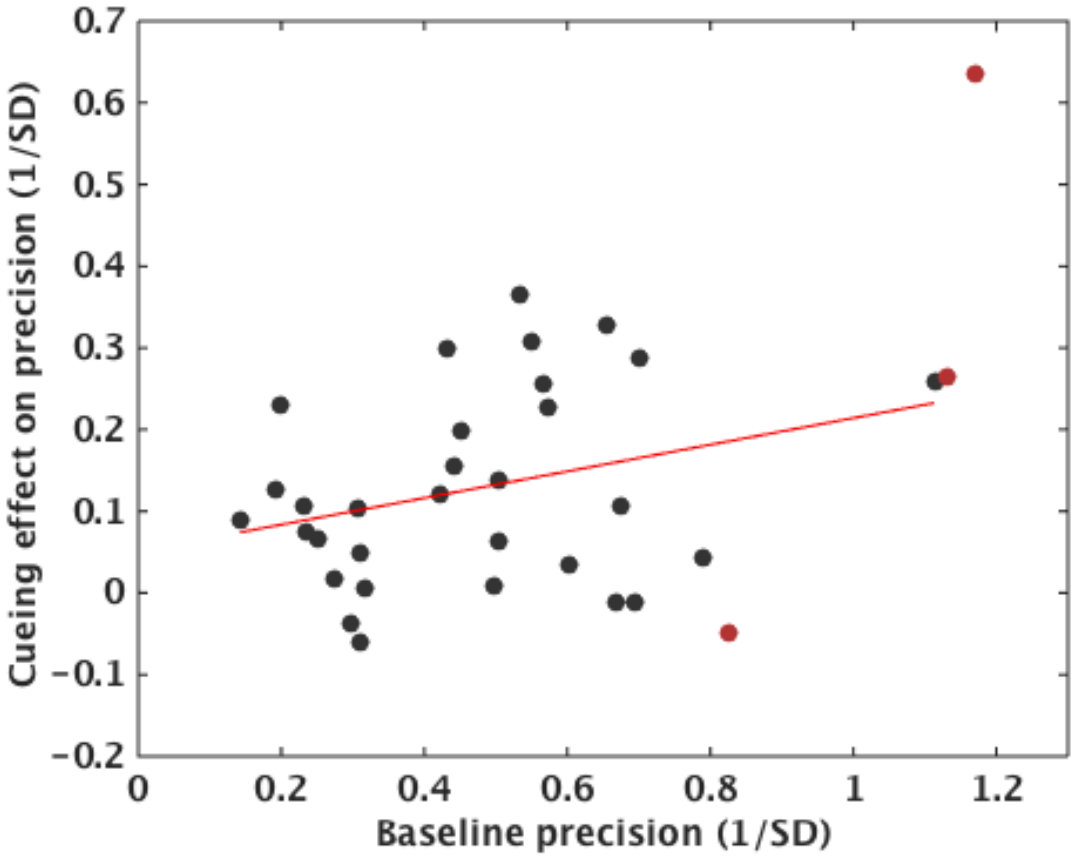

Figure 4-14. Baseline performance (precision, averaged over contrast levels 3%, 10%, and 50% in the neutral retrocue condition) had a tendency to predict magnitude of the retrocueing effect on recall precision, p = 0.108. Outliers (N = 3, shown in red) were removed when they had a Cook's distance statistic greater than three times the mean Cook's distance. Without the exclusion of outliers the regression equation was significant with an R^2 of .202 (adjusted R^2 = .177), $F(1,32) = 8.103$, p = .008, β = .450.

4.5.8 Interim Discussion

As in Experiments 1 and 2, the results showed that recall precision is better following valid compared to neutral retrocues. This effect was moderated by individual differences in baseline recall precision. Median-split analysis revealed stronger retrocueing benefits for high performers compared to low performers. Similarly, linear regression analysis revealed a strong trend for a positive relationship between baseline precision and retrocueing magnitude in recall precision. A novel finding is that for the

contrast levels used in this study (3, 10 and 18%) retrocueing effects were modulated not only by individual differences, but also by contrast level. Specifically, retrocueing effects on recall precision increased from 3% contrast to 10% contrast and then remained stable. This suggests that Experiments 1 and 2 may have missed the elbow of the psychometric function, and that the representational quality of WM items is modulated by contrast levels.

Applying a mixture model to the data allowed me to examine the pattern of results more closely and I observed that kappa was not modulated by contrast or retrocueing. This is in line with the low-contrast condition of Experiment 2 where no effects of retrocueing were observed on kappa. Together, this suggests that retrocues only affect kappa, i.e. the representational quality of WM items, at high levels of contrast. Furthermore, I now observed an effect of retrocueing on misbinding rates and misbinding decreased following valid compared to neutral retrocues. Thus, compared to Experiment 2, the effects of valid retrocueing on misbinding and kappa reversed and retrocueing benefits were driven by a reduction in misbinding rates (in addition to increased target recall and decreased guess rates). Overall misbinding rates in neutral retrocueing trials were higher in Experiment 3 compared to Experiment 2 (although not significantly, M_{E2} = 4% ±0.8%, M_{E3} = 5.8% ±0.7%; p = 0.103). Possibly, misbinding was unaffected by retrocueing in Experiment 2 because misbinding rates were already low in neutral retrocueing conditions. Since task conditions and overall performance were identical except for the change in contrast, it appears that decreasing the quality of incoming

sensory information leads to an increase in misbinding errors, which then allows for retrocueing benefits.

The effects on target response probability and guess rate were comparable to Experiment 2 and neither one was affected by individual differences in baseline performance or by contrast. Target-response probability was increased following valid retrocues, as in Experiment 2. Guess rate was not affected by retrocueing or by contrast whereas I observed a strong trend for retrocueing in Experiment 2.

4.2. Discussion

The results indicate that the quality of incoming sensory information interacts with attentional orienting to shape the representational quality of items in WM. This effect is strongly modulated by individual differences in baseline task performance and together these variables determine how incoming sensory information is encoded into WM. Additionally, retrocueing effects on overall recall precision were modulated by the quality of incoming sensory information and were more pronounced at higher levels of contrast, specifically at 10%+ contrast compared to 3% contrast. Using a model-based approach I found evidence that the process by which retrocueing affects WM recall is dependent on the quality of incoming sensory information. While effects of retrocueing on target recall and guess rate were relatively constant across Experiment 2 and 3, and across different contrast levels, effects of retrocueing on misbinding and kappa differed between experiments and contrast levels. At high levels of contrast (Experiment 2) effects of retrocueing were mainly observed on kappa, while at low levels of contrast (Experiment 3) retrocueing effects were observed

on misbinding rates (in addition to affecting guess rate equally at every contrast).

The relative proportion of misbinding often varies between studies ranging from no misbinding up to 30% (Matthey et al., 2015). Matthey et al. (2015) recently proposed a formal account of misbinding errors. According to this model, representations for different WM items overlap because one population of neurons encodes all items. WM performance is limited by neural noise and item-overlap, and the amount of overlap is determined by how items are encoded. Overlapping items can lead to misbinding errors as items interfere with one another. The balance between two kinds of units, which either store individual features or binding information, determines the relative proportion of misbinding errors and the representational quality of items in WM (kappa). However, this model does not address how changes in noise modulate misbinding versus kappa. A model proposed by Bays (2014) attributes WM recall errors to the presence of noise in the neural populations that code WM items. Noise limits the decodability of neural populations and item-related activity diffuses over time. Increases in noise can be countered by increases in gain, which amplifies the strength of the coding signal. Although this model only considers internal but not external sources of noise, it does provide a framework to consider the effects of external sources of noise on representational quality of WM items.

Combining these two models allows to offer tentative explanations on how noise fluctuations and attentional mechanisms interact to shape the representational quality of and relative misbinding rates between items in WM. My results support the suggestion that overlap and interference between items in WM depends on how they are encoded (Matthey et al.,

2015). When the quality of sensory input is low, such as when items are presented at low contrast, items are encoded with more overlap and interference between items increases. Attentional selection (e.g., from retrocueing) then increases the gain and signal-to-noise ratio (SNR) of individual representations, separating noisy, overlapping items. This allows for categorical distinctions to be drawn between representations, reducing the potential for misbinding errors. At higher contrast levels with a higher SNR, different items are encoded with less overlap and attentional selection does not affect item-overlap as items are already (relatively) readily distinguished – at least in the experimental paradigm employed here. Attentional selection and increases in gain do not further differentiate item-representations but instead the increase in gain sharpens item representations, thereby increasing the quality of memory recall (as measured by the concentration parameter, kappa). These effects were additionally modulated by individual differences in baseline performance. Specifically, a median-split analysis revealed group by retrocue interactions with significantly higher retrocueing effects on recall precision in high-performers. I also observed a strong positive relationship between recall precision in neutral retrocueing conditions and the magnitude of the retrocueing effect on recall precision. Individuals with high baseline recall precision benefitted more from valid retrocueing than worse-performing individuals. This positive relationship was observed in all three experiments, although in Experiment 3 it only approached significance. Future studies are needed to disentangle the relative influences of the quality of incoming sensory information, attentional selection and individual differences in contrast sensitivity and WM capacity on misbinding rates.

In addition, I did not find evidence that items presented below a level of conscious awareness can be ‘boosted’ into a conscious WM representation, in contrast to Sergent et al.'s (2013) findings. A follow-up study by the same group (Thibault et al., 2016) that was published after completion of the Experiments presented here showed similar results with objective measures of task performance. They used a task procedure similar to Sergent et al. (2013), but assessed WM recall performance rather than subjective ratings. Participants were asked to report back the exact stimulus orientation using a continuous scale and the stimulus contrast was staircased so responses always fell within 45° of the target value (at ~3% Michelson contrast). They found that valid retrocues reduced the guess rate but did not affect the quality of WM representation, consistent with their previous proposal that retrospective attention can render invisible stimuli visible. Together, these studies contrast both the effects on subjective ratings and the effects on WM recall I report in Experiment 1.

There are two key differences between the studies reported in this Chapter and the design used by Sergent et al. (2013) and Thibault et al. (2016). Regarding the subjective ratings, in my task observers were instructed to rate the visibility of all three stimuli during initial stimulus presentation while retrospective attention was only directed at one of those three stimuli. In Sergent et al. (2013) there was only one stimulus. Even if observers’ subjective ratings of the stimulus visibility were changed by subsequent attentional selection, this would have affected only one out of three stimuli while ratings were made on all three stimuli. Additionally, in my task subjective ratings were made after participants had received feedback regarding their performance and this is likely to have biased their ratings.

Second, both Sergent et al. (2013) and Thibault et al. (2016) presented cues at varying time points before or after target presentation, between 400 ms before and 400 ms after target presentation (± 350 ms in the case of Thibault et al., 2016). Cues in this study were presented 1,000 ms after stimulus offset. By the time retrospective attentional selection occurred, any item-related activity may already have decayed. Both of these variables could account for the differences in results and should be examined in future studies.

5 General Discussion

The aim of this thesis was to establish a fuller and more nuanced description of biases in information processing. I used behavioural, MEG, and fMRI measures to contrast feature-based and spatial attention mechanisms, characterize interactions between perceptual and memory presentations, and examine retrospective attentional mechanisms within WM. The specific results from each study have already been discussed in the respective chapters. Therefore, this chapter will focus on discussing the results in their broader context, considering their limitations, and suggesting future directions of research. I will only briefly summarize the results. The following discussion is divided according to the main themes of the thesis.

5.1 Similarities and differences between feature-based and spatial attention

In Chapter 2 of this thesis, I examined the degree to which similar neural mechanisms underlie feature-based attention and spatial attention. I recorded brain activity using fMRI and MEG, in two separate sessions, while participants completed a target-detection task that independently manipulated feature-based and spatial attention. Behaviourally, feature-based and spatial attention independently improved performance and there was no evidence of interactive effects on performance. The fMRI results showed that largely identical regions were activated following spatial cues and feature cues, a broad network of task-related activity consistent with previous reports of the frontoparietal control network (e.g. Shulmann et al., 2002; Yantis et al.2002). The MEG results provided a much more temporally

fine-tuned description of attention-related brain activity. For analysis of the MEG data, I focused on oscillatory components in the alpha and gamma frequency band, consistent with an overwhelming literature reporting contributions of these brain rhythms to attentional mechanisms. I will discuss the alpha-band and gamma-band results in turn.

The findings in the alpha band reveal that alpha power modulations are under flexible top-down control and can bias neural processing on the basis of feature and spatial information, depending on the current task requirements. Alpha lateralization during the cue target interval only somewhat differentiated between feature-based and spatial preparatory attention: comparison of the alpha lateralisation following spatial cues with feature information to spatial cues without feature information revealed a positive cluster in the cue-target interval from 440 ms to 620 ms after cue onset. This difference appeared to reflect a difference in the timing of alpha lateralization following the different types of spatial cues: timecourses of the alpha lateralization revealed a relative delay in the alpha modulation when spatial cues did contain feature information compared to when they did not. Further analysis focusing on non-lateralised alpha power reliably differentiated between the two cue types. First, non-lateralised alpha power was relatively more suppressed following feature cues (i.e. feature-spatial and feature cues) compared to non-feature cues (i.e. neutral and spatial cues). This difference extended bilaterally, but was more pronounced over left channels. To assess the spatial specificity of the effect, I compared non-lateralised alpha power following left (right) spatial cues with versus without feature information. This revealed that the reduction in alpha power

following feature cueing occurred primarily over ipsilateral channels: alpha suppression was greater over left (right) channels for left (right) spatial cues with feature information compared to without feature information. Alpha power over contralateral sensors was also reduced by the presence of feature information in spatial cues, but this was less pronounced than over ipsilateral sensors. In other words, feature-spatial cues led to relatively less pronounced alpha lateralization than spatial cues. Instead, alpha suppression occurred over both ipsilateral and contralateral channels.

I argue that this reflects a relatively more global modulation of alpha power following feature-based attention. Most likely, the effect of feature information on contralateral alpha power is relatively small as the dynamic range of alpha power is limited, and alpha power can only be suppressed to a certain extent. Most of the effect of feature information, therefore, will be observed over ipsilateral channels where alpha power remains relatively high with just spatial cues. In short, adding feature information to spatial cues reduces the difference between contralateral and ipsilateral alpha power. An analysis approach focusing on alpha lateralization is essentially blind to the differences between spatial cues with and without feature information. Analysis of non-lateralised alpha power, on the other hand, revealed that both feature-based attention modulates alpha power in a more global fashion than spatial attention, a finding that is line with the typical finding of spatially distributed effects of feature-based attention on behaviour and BOLD responses (Folk et al., 2002; Saenz et al., 2002; Serences et al., 2005).

It is conceivable that the more pronounced alpha power suppression following feature cues is related to differences in task difficulty (Brouwer, Hogervorst, Holewijn, & van Erp, 2014; Fairclough, Venables, & Tattersall, 2005; Klimesch, Schimke, & Pfurtscheller, 1993). Perhaps preparing for the upcoming target-detection event is more taxing when observers are given two pieces of information, i.e. feature and spatial information, compared to conditions when only spatial information is used for attentional guidance. Alternatively, the use of feature information, compared to spatial information, might be more difficult. Either way, a difference in difficulty might have lead to the observed increase in alpha power suppression. There are several reasons why this explanation is unlikely to account fully for the feature-related alpha suppression I observed. First, at the behavioural level performance was comparable between valid spatial and valid feature cueing conditions. Secondly, the neutral cueing condition is by far the most difficult, as reflected in RTs and accuracy, yet the comparison of feature cues to non-feature cues (which contained spatial and neutral cues) revealed greater alpha power suppression following feature cues. Finally, the spatial specificity of the additional alpha power suppression is also evidence against an interpretation whereby overall differences in task difficulty drove the differences in alpha modulations.

The additional modulation of alpha power by feature information did not persist throughout the entire cue-target interval and was only observed between ~420 ms to ~800 ms after cue onset. As already mentioned, perhaps this effect merely reflects differences in cue processing and is not related to feature-based preparatory attention mechanisms. However, it is

unlikely that differences in cue processing would persist until ~800 ms after cue onset (600 ms after cue offset).

As an alternative explanation, the short-lived nature of this difference might be related to the proposed functional roles of alpha power modulation. The functional inhibition account (Klimesch et al., 2007) proposes that alpha modulation supports two processes. A decrease in alpha power reflects a release from inhibition of neurons coding for task-relevant information, while an increase reflects active inhibition of neurons coding for task-irrelevant information. Possibly, the relative contributions of the two mechanisms changed over the course of the cue-target interval in the experiment reported here. I have argued that feature information leads to additional alpha power suppression of task-relevant neurons over ipsilateral channels. However, because neurons with different feature-selectivity profiles overlap, additional alpha power suppression would only be evident in the average ipsilateral data if task-irrelevant neurons over ipsilateral channels do not show alpha power enhancement. If both mechanisms occur, the effects cancel each other out and no additional alpha power suppression is observed over ipsilateral neurons. Perhaps, the short-lived nature of the enhanced alpha power decrease over ipsilateral channels reflects additional active suppression of task-irrelevant neurons at later points during the cue-target interval. Future studies could tease apart the relative contributions of active inhibition versus release from inhibition throughout preparatory intervals.

The findings in the gamma band suggest a special role of gamma power in feature-based attention. Specifically, gamma power increased

contralateral to the side of target presentation only following feature cues. I did not observe such target-induced gamma lateralization in any of the other conditions. In other words, there was a significant target-induced gamma response only when observers were maintaining a feature-specific target template as a target was presented. This is consistent with the Match and Utilization Model (MUM, Herrmann et al. 2004) according to which early gamma responses reflect the match between top-down expectations and bottom-up stimulus input. Specifically, when observers are maintaining a feature-specific target template, top-down influences induce subthreshold oscillations in target populations. Incoming signals that match then enhance the existing signal. According to the MUM, the target-induced gamma response should be largest when observers maintain both feature-specific and location-specific templates regarding upcoming stimulus input. I did not formally examine this[9], however, an exploratory comparison of the gamma lateralisation indices for the feature vs nonfeature comparison between cues with spatial information and cues without spatial information revealed no differences. Future studies should assess the specific contributions of feature and spatial expectations on the early, target-related gamma response.

5.2 Interfering effects of perceptual information on WM representations

In Chapter 3, I asked how exposure to supraliminal and subliminal distracter stimulus, before or after target encoding, affects WM

[9] Doing this comparison would require splitting up the data by an additional factor (presence or absence of spatial information in the cue) which halves the trial numbers per condition, substantially reducing experimental power.

representations. Does the presentation of a distracter degrade target-related information and/or does it alter the information (and if so, how)? I found that WM recall was influenced by the presentation of supraliminal distracter items. This influence was expressed in two ways. First, WM recall was shifted toward the distracter value. Second, WM recall precision worsened with distracter presentation. These effects were robust and emerged both when the distracter appeared before, and when it appeared after, the target item. I also contrasted the effects of supraliminal distracters with subliminal distracters. When distracters were presented below a level of conscious awareness, WM recall performance was no longer influenced by the presentation of a distracter. Indeed, conditions with subliminal distracters were statistically indistinguishable from conditions without distracters. Finally, I manipulated the task-relevance of the distracter item to assess how task-relevance modulates effects of the supraliminal or subliminal distracters. I did not find evidence that manipulations of task-relevance or selective attention modulated the effects of subliminal or supraliminal distracters.

These findings provide a basis to assess the predictions made from two accounts of distracter interference effects in WM. Magnussen (2000) explains distracter interference effects by proposing suppressive interactions between feature-specific memory stores. If two stores contain conflicting information, lateral inhibition between the stores will degrade the stored information. This can explain the findings from various studies that have shown interference effects as long as distracters are not identical to target items. The perceptual averaging account (Dubé et al., 2014; Huang &

Sekuler, 2010), on the other hand, proposes that distracter interference effects arise from an averaging process that establishes a target representation by averaging over target and distracter values according to their assigned weights, with weights determined by selective attention mechanisms.

The feature-specific memory stores account cannot (Magnussen, 2000) fully explain the set of findings reported here, as I reported not just information loss, but also a change of information, following distracter presentation. However, the perceptual averaging account (Dubé et al., 2014; Huang & Sekuler, 2010) also makes predictions that are at odds with the findings presented here. First, I found that the biasing effect was strongest when distracters were about 40° different to the target value, and then weakened as the difference between target and distracter values increased. According to the perceptual averaging account, however, distracter interference effects should increase the more dissimilar distracters are from targets.

Furthermore, I did not find evidence that manipulations of task-relevance or selective attention modulated the effects of subliminal or supraliminal distracters. Again, this is different to the predictions made by the perceptual averaging account and is also different to the findings reported by Huang and Sekuler (2010) and Dubé et al. (2014). In both of these studies, distracter effects were substantially reduced when predictive cues were presented before the memory items. Possibly, these differences could be explained by differences in task design: in the studies presented in this thesis participants were never presented with pre-cues. Instead, the key

difference was that in Experiment 2 the information identifying the target item was presented immediately after presentation of the first WM item (i.e. if a mask was presented after the first item, the second item was the target, if there was no mask then this item was the target), whereas in Experiment 3 cues were presented only after both WM items had been presented. This design does not optimize task-relevance/selective attention differences as there is some ambiguity regarding the target identity even in Experiment 2. Another possible explanation for the differences might be that I explicitly modelled misbinding, whereas Huang and Sekuler (2010) and Dubé et al. (2014) did not. Wallis et al. (2015) showed that precues reduced misbinding rates relative to neutral cueing conditions. A reduction in misbinding rates would lead to an apparent reduction of the biasing effect as misbinding trials contribute to the overall magnitude of the bias.

Finally, the perceptual averaging account clearly predicts the direction of the distracter bias: distracter presentation should always bias target recall toward the distracter value. While I observed this pattern at the group level, as reported in Chapter 3, there were pronounced individual differences not only in the magnitude but also in the direction of the effect. In both Experiments, for some participants target recall was biased *away* from the distracter value. The perceptual averaging account in its present form cannot accommodate this finding.

A recent study published after completion of the Experiments presented in Chapter 3 reports a number of relevant findings (Rademaker, Bloem, De Weerd, & Sack, 2015). Similar to the WM task in Chapter 3, observers completed a sequential WM task for two items and WM recall

precision and bias were measured. In one version of the task observers were explicitly instructed to ignore the second item. In another version, a postcue indicated the task-relevant item. Rademaker et al. (2015) found that the bias effect was stronger when distracters had to be ignored, which is the opposite of the predictions made by the perceptual averaging account. Furthermore, the authors also reported a reversal of the direction of the bias effect in some conditions. When observers were cued to report the first item, recall bias was in the direction of the distracter value. However, when the second item was probed, recall was biased away from the distracter value in five out of eight participants. This directly contradicts findings by Huang and Sekuler (2010) where memory reports were biased toward the distracter irrespective of what order the items were presented in.

Together, the results I present in this thesis and the results discussed, paint a complicated picture. In Rademaker et al. (2015) task-irrelevant distracter exerted a stronger biasing effect, in Huang and Sekuler, (2010) and Dubé et al. (2014) task-irrelevant distracters exerted weaker biasing effects, and I did not observe any differences in recall bias between different conditions of task-relevance. Similarly, both my work and Rademaker et al. (2015) reports recall biases toward as well as away from the distracter value. Clearly, the perceptual averaging account cannot fully explain these sets of findings.

What then might be a complete explanation of the varied pattern of results found in comparable tasks? I have already described the difficulties in interpreting results that do not account for misbinding errors. It is difficult to debate the nature of WM representations when it is unclear whether

reported findings truly reflect effects on the WM representation of interest. Future studies need to design task conditions carefully to untangle the nature of relationship between existing memory representations and incoming sensory information.

Both Rademaker et al. (2015) and I presented evidence that the distracter biasing effect disappears when the distracter is rendered subliminal, findings that directly contrast with reports from Silvanto and Soto (2012). In Rademaker et al. (2015) stimuli were rendered subliminal by presenting them dichoptically to induce binocular rivalry, whereas I used pattern masks as in Silvanto and Soto (2012). As already discussed in Chapter 3, the differences between the results reported by Silvanto and Soto (2012) and the results presented here, are most likely related to differences in how awareness was assessed in the two studies. The fact that Rademaker et al. (2015) and I reported similar findings while relying on different masking techniques, which potentially interfere at different stages of visual processing, suggests that, at most, subliminal distracters could only exert a limited influence on WM processes, and WM representations are protected from subliminal influences.

5.3 Interactions within WM depend on sensory signal strength during encoding

In Chapter 4, I asked how WM performance benefits from retrospective attention. Specifically, I examined whether the nature of the retrocueing benefit relates to the representational quality of items encoded into WM. To assess the influence of changes in the representational quality of items in WM on the retrocueing benefit, I varied stimulus contrast levels

during encoding, and considered individual differences in baseline task performance. The results reported in Chapter 4 show that retrocueing mechanisms are flexible, and retrocues reduce different sources of error depending on the quality of WM representations. Valid retrocues reduced guess rates and increased the probability of correct target recall at all contrast levels. However, at high contrast levels, valid retrocues additionally improved recall precision, increasing the fidelity of individual WM representations. At low contrast levels, valid retrocues reduced misbinding rates, targeting interactions between items in WM rather than the fidelity of individual item representations. This pattern was further modulated by individual differences in baseline WM performance, with differences in cueing effects that were induced by changes in contrast level further exacerbated by baseline levels of performance. This suggests that the mechanisms underlying the retrocueing benefit are flexible, and target different aspects of WM depending on the representational quality of items within WM.

The results presented in Chapter 4 show that changes in contrast levels during encoding qualitatively change top-down attentional effects in WM. Interactions of attentional mechanisms with bottom-up stimulus properties have been demonstrated previously. In perception, different models have characterized the exact nature of the interaction between contrast levels and attention. For one, studies have shown that attention appears to shift neurons' contrast response function to the left to increase the effective contrast of the stimulus in the neurons' receptive field (Martínez-Trujillo & Treue, 2002; Reynolds et al., 2000). These results are

consistent with the contrast-gain model of attention, which proposes that effects of attention are comparable to increases in stimulus contrast and lead to a leftward shift of the contrast-response function. However, other studies using neurophysiological, neuroimaging, and behavioural methods, have provided conflicting results that are more in line with the response gain model of attention, which predicts that attention increases firing rate by a multiplicative response gain factor (Prinzmetal et al., 1997; Liston and Stone, 2008; Schneider and Komlos, 2008; Treue and Maunsell, 1996; Treue and Trujillo, 1999; McAdams and Maunsell, 2000; Williford and Maunsell, 2006). Neither the contrast gain model, which predicts that effects of attention are strongest at low contrast levels, nor the response gain model, which predicts that effects of attention are strongest at high levels of contrast, can fully capture the empirical evidence in the field. The emerging picture from these studies is that attention and contrast interact in different ways within perception depending on stimulus properties and the stage of visual processing (Huang & Dobkins, 2005; Reynolds & Heeger, 2009). The findings presented in Chapter 4 suggest that in WM, like in perception, the effects of attention can vary and depend on the quality of internal representations, as determined by bottom-up stimulus properties. More research is needed for a detailed understanding of the mechanisms underlying the observed effects on recall precision and misbinding.

Changes in the stimulus contrast level during encoding had effects on baseline performance (in neutral cueing conditions), consistent with the argument made here that differences in bottom-up stimulus properties influence the fidelity of internal representations, accounting for some of the

variability observed in WM performance across items and trials. However, the effect of stimulus contrast on WM performance was only modest once stimuli were presented at a level that allowed for above-chance performance. There was a slight increase in baseline performance from 3% contrast to 10% contrast but from then on performance was comparable between different contrast levels. As there are individual differences in contrast sensitivity (Baker, 2013), matching stimuli contrast levels to subjects' individual detection thresholds might have helped to establish parametric changes in baseline WM performance. However, I chose not to do this because it is not straightforward to translate results from standard forced-choice staircasing tasks to the WM task used in Chapter 4. In future studies, contrast levels could be titrated based on performance in the WM task, for example by varying the recall precision necessary for a response to be classified as correct (see for example Thibault et al., 2016).

I also observed large individual differences in WM performance. These might be related to individual difference in WM ability, or individual differences in task strategy. For example, a series of studies have examined whether the trade-off between WM quality and quantity is under voluntary control (Fougnie, Cormiea, Kanabar, & Alvarez, 2016; Murray, Nobre, Astle, & Stokes, 2012; Zhang & Luck, 2011). These studies have provided mixed results. Zhang and Luck (2011) and Murray et al. (2012) did not find evidence that participants could voluntarily trade-off quality and quantity in WM. Murray et al. (2012) manipulated task requirements and participants' task expectations across four experiments to examine whether participants traded-off quality and quantity in line with task demands, and found no

evidence for this. For example, recall precision for a stimulus array did not increase when greater precision was required to complete the task compared to when only low precision was required. Similar results were found when expected set size (and not expected precision to complete the task) was manipulated, as well as when both expectations were combined to maximize effects. Fougnie et al. (2016), however, did find evidence that participants can trade-off quantity and quality along with task instructions. Here, instructions varied to encourage participants either to maximize recall precision or the number of items remembered for a display of five colour disks. Specifically, in separate blocks participants were told they would earn bonus points, which translated to financial rewards, if a) they minimized recall error for one cued item or b) they reported all five items within 90° of the correct value. Both recall precision and guess rate were reduced when all items had to be reported at a relatively coarse precision compared to when only one item had to be reported at high precision, consistent with the proposal that quality and quantity of WM items are under flexible, top-down control. This pattern persisted even when task requirements varied randomly throughout the experiment, and trial types were not revealed until after encoding had taken place; but disappeared when trial type was revealed at the response stage.

It is not clear if, or to what degree, changes in task strategy contributed to the changes in the retrocueing benefit across individuals in the experiments reported in Chapter 4. However, the fact that top-down trade-offs between quality and quantity were only reported in tasks with relatively strong incentives (in the form of real-world financial rewards)

suggests that such trade-offs are unlikely to have substantially contributed to my findings. Instead, it is more likely that individual differences in WM ability accounted for variability in the retrocueing benefit across individuals. This, in combination with the findings that changes in cueing effects at different contrast level were enhanced, and not qualitatively changed, by differences in baseline performance levels, suggests that both contrast and baseline WM performance levels changed the representational quality of items in WM, which modulated effects of retrospective attentional selection. Future studies could more directly test this by recruiting larger sample sizes and assessing whether estimates of individuals' WM abilities (e.g. WM capacity using Cowan's K; Cowan, 2001) can predict the effects of retrospective attentional selection.

5.4 Conclusions

In this thesis, I described a series of experiments investigating the various ways in which information processing can be biased to shape perception and working memory. In Chapter 2, I demonstrate that alpha power modulations are under flexible top-down control and can bias information processing in favour of whichever information is currently task-relevant. The results reported in Chapter 3 show that WM representations, unlike perceptual representations, are not susceptible to subliminal influences. However, supraliminal information can subtly bias WM representations demonstrating a permeability of internal representations long after encoding into WM has taken place. The work presented in Chapter 4 extends the research on retrospective attention within WM to demonstrate that different mechanisms produce the retrocueing benefit, and

retrospective attention flexibly adapts to the representational quality of items in WM. Together, the findings highlight how information processing is biased at multiple stages, that these biases have different origins, and manifest themselves in a variety of ways.

6 References

Adam, K. C. S., Mance, I., Fukuda, K., & Vogel, E. K. (2015). The Contribution of Attentional Lapses to Individual Differences in Visual Working Memory Capacity. *Journal of Cognitive Neuroscience*, *27*(8), 1601–1616. http://doi.org/10.1162/jocn_a_00811

Adjamian, P., Holliday, I. E., Barnes, G. R., Hillebrand, A., Hadjipapas, A., & Singh, K. D. (2004). Induced visual illusions and gamma oscillations in human primary visual cortex. *European Journal of Neuroscience*, *20*(2), 587–592. http://doi.org/10.1111/j.1460-9568.2004.03495.x

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. http://doi.org/10.1177/0956797610363543

Ansorge, U., & Neumann, O. (2005). Intentions Determine the Effect of Invisible Metacontrast-Masked Primes: Evidence for Top-Down Contingencies in a Peripheral Cuing Task. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(4), 762–777. http://doi.org/10.1037/0096-1523.31.4.762

Arnal, L. H., & Giraud, A.L. (2012). Cortical oscillations and sensory predictions. *Trends in Cognitive Sciences*, *16*(7), 390–398.

Arnal, L. H., Wyart, V., & Giraud, A.L. (2011). Transitions in neural oscillations reflect prediction errors generated in audiovisual speech. *Nature Neuroscience*, *14*(6), 797–801. http://doi.org/10.1038/nn.2810

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

Axmacher, N., Cohen, M. X., Fell, J., Haupt, S., Dümpelmann, M., Elger, C. E., et al. (2010). Intracranial EEG Correlates of Expectancy and Memory Formation in the Human Hippocampus and Nucleus Accumbens. *Neuron*, *65*(4), 541–549. http://doi.org/10.1016/j.neuron.2010.02.006

Baars, B. J. (1988). A cognitive theory of consciousness, Cambridge University Press, Cambridge.

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

Bae, G.Y., Olkkonen, M., Allred, S. R., Wilson, C., & Flombaum, J. I. (2014).

Stimulus-specific variability in color working memory with delayed estimation. *Journal of Vision*, *14*(4), 7–7. http://doi.org/10.1167/14.4.7

Baker, D. H. (2013). What is the primary cause of individual differences in contrast sensitivity? *PLoS ONE*, *8*(7), e69536. http://doi.org/10.1371/journal.pone.0069536

Baldassi, S., & Verghese, P. (2005). Attention to locations and features: different top-down modulation of detector weights. *Journal of Vision*, *5*(6), 556–570. http://doi.org/10.1167/5.6.7

Baldauf, D., & Desimone, R. (2014). Neural mechanisms of object-based attention. *Science*, *344*(6182), 424–427. http://doi.org/10.1126/science.1247003

Bauer, M., Kennett, S., & Driver, J. (2012). Attentional selection of location and modality in vision and touch modulates low-frequency activity in associated sensory cortices. *Journal of Neurophysiology*, *107*(9), 2342–2351. http://doi.org/10.1152/jn.00973.2011

Bauer, M., Oostenveld, R., Peeters, M., & Fries, P. (2006). Tactile spatial attention enhances gamma-band activity in somatosensory cortex and reduces low-frequency activity in parieto-occipital areas. *Journal of Neuroscience*, *26*(2), 490–501. http://doi.org/10.1523/JNEUROSCI.5228-04.2006

Bauer, M., Stenner, M.-P., Friston, K. J., & Dolan, R. J. (2014a). Attentional modulation of alpha/beta and gamma oscillations reflect functionally distinct processes. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *34*(48), 16117–16125. http://doi.org/10.1523/JNEUROSCI.3474-13.2014

Bauer, M., Stenner, M.-P., Friston, K. J., & Dolan, R. J. (2014b). Attentional modulation of alpha/beta and gamma oscillations reflect functionally distinct processes. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *34*(48), 16117–16125. http://doi.org/10.1523/JNEUROSCI.3474-13.2014

Baylis, G. C., & Driver, J. (1992). Visual parsing and response competition: The effect of grouping factors. *Perception & Psychophysics*, *51*(2), 145–162. http://doi.org/10.3758/BF03212239

Baylis, G. C., & Driver, J. (1993). Visual attention and objects: evidence for hierarchical coding of location. *Journal of Experimental Psychology: Human Perception and Performance*, *19*(3), 451–470.

Bays, P. M. (2014). Noise in neural populations accounts for errors in working memory. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *34*(10), 3632–3645. http://doi.org/10.1523/JNEUROSCI.3204-13.2014

References

Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, *321*(5890), 851–854. http://doi.org/10.1126/science.1158023

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. http://doi.org/10.1167/9.10.7

Bays, P. M., Gorgoraptis, N., Wee, N., Marshall, L., & Husain, M. (2011). Temporal dynamics of encoding, storage, and reallocation of visual working memory. *Journal of Vision*, *11*(10), 6–6. http://doi.org/10.1167/11.10.6

Behrmann, M., & Tipper, S. P. (1999). Attention accesses multiple reference frames: evidence from visual neglect. *Journal of Experimental Psychology: Human Perception and Performance*, *25*(1), 83–101. http://doi.org/10.1037/0096-1523.25.1.83

Bennett, P. J., & Cortese, F. (1996). Masking of spatial frequency in visual memory depends on distal, not retinal, frequency. *Vision Research*, *36*(2), 233–238. http://doi.org/10.1016/0042-6989(95)00085-E

Berens, P. (2009). CircStat: a MATLAB toolbox for circular statistics. *Journal of Statistical Software*.

Bichot, N. P., Heard, M. T., DeGennaro, E. M., & Desimone, R. (2015). A Source for Feature-Based Attention in the Prefrontal Cortex. *Neuron*, *88*(4), 832–844. http://doi.org/10.1016/j.neuron.2015.10.001

Bichot, N. P., Rossi, A. F., & Desimone, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area V4. *Science*, *308*(5721), 529–534. http://doi.org/10.1126/science.1109676

Bollimunta, A., Mo, J., Schroeder, C. E., & Ding, M. (2011). Neuronal mechanisms and attentional modulation of corticothalamic α oscillations. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *31*(13), 4935–4943. http://doi.org/10.1523/JNEUROSCI.5580-10.2011

Bona, S., Cattaneo, Z., Vecchi, T., & Soto, D. (2013). Metacognition of visual short-term memory: dissociation between objective and subjective components of VSTM. *Frontiers in Psychology, 4.*

Brady, T. F., & Alvarez, G. A. (2010). Ensemble statistics of a display influence the representation of items in visual working memory. Visual Cognition.

Brady, T. F., & Alvarez, G. A. (2011). Hierarchical encoding in visual working memory: ensemble statistics bias memory for individual items.

Psychological Science, *22*(3), 384–392. http://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. *Journal of Vision*, *11*(5), 4–4. http://doi.org/10.1167/11.5.4

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436.

Broadbent, D. (1958). Perception and Communication (Pergammon, Oxford).

Broadbent, D. E. (1982). Task combination and selective intake of information. *Actpsy*, *50*(3), 253–290.

Brouwer, A.M., Hogervorst, M. A., Holewijn, M., & van Erp, J. B. F. (2014). Evidence for effects of task difficulty but not learning on neurophysiological variables associated with effort. *International Journal of Psychophysiology*, *93*(2), 242–252. http://doi.org/10.1016/j.ijpsycho.2014.05.004

Buffalo, E. A., Fries, P., Landman, R., Buschman, T. J., & Desimone, R. (2011). Laminar differences in gamma and alpha coherence in the ventral stream. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(27), 11262–11267. http://doi.org/10.1073/pnas.1011284108

Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, *97*(4), 523–547. http://doi.org/10.1037/0033-295X.97.4.523

Carrasco, M. (2014). Spatial covert attention: Perceptual modulation. *The Oxford Handbook of Attention*.

Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience*, *7*(3), 308–313. http://doi.org/10.1038/nn1194

Cohen, M. R., & Maunsell, J. H. R. (2010). A neuronal population measure of attention predicts behavioral performance on individual trials. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *30*(45), 15241–15253. http://doi.org/10.1523/JNEUROSCI.2171-10.2010

Cohen, M. R., & Maunsell, J. H. R. (2011). Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron*, *70*(6), 1192–1204. http://doi.org/10.1016/j.neuron.2011.04.029

Corbetta, M., Kincade, J. M., Ollinger, J. M., McAvoy, M. P., & Shulman, G. L. (2000). Voluntary orienting is dissociated from target detection in

human posterior parietal cortex. *Nature Neuroscience*, *3*(3), 292–297. http://doi.org/10.1038/73009

Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1990). Attentional modulation of neural processing of shape, color, and velocity in humans. *Science*, *248*(4962), 1556–1559.

Corbetta, M., Miezin, F. M., Shulman, G. L., & Petersen, S. E. (1993). A PET study of visuospatial attention. *Journal of Neuroscience*, *13*(3), 1202–1226.

Cornelissen, F. W., Peters, E. M., & Palmer, J. (2002). The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior Research Methods, Instruments, & Computers : a Journal of the Psychonomic Society, Inc*, *34*(4), 613–617.

Cowan, N. (1995). Attention and memory: an integrated framework. http://doi.org/10.1300/J025v06n03_10

Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*(1), 87–114– discussion 114–85.

Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, *7*(9), 415–423. http://doi.org/10.1016/S1364-6613(03)00197-9

David, S. V., Hayden, B. Y., Mazer, J. A., & Gallant, J. L. (2008a). Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. *Neuron*, *59*(3), 509–521. http://doi.org/10.1016/j.neuron.2008.07.001

David, S. V., Hayden, B. Y., Mazer, J. A., & Gallant, J. L. (2008b). Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. *Neuron*, *59*(3), 509–521. http://doi.org/10.1016/j.neuron.2008.07.001

Dehaene, S., Naccache, L., Le Clec'H, G., Koechlin, E., Mueller, M., Dehaene-Lambertz, G., et al. (1998). Imaging unconscious semantic priming. *Nature*, *395*(6702), 597–600. http://doi.org/10.1038/26967

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*(1), 193–222. http://doi.org/10.1146/annurev.ne.18.030195.001205

Deutsch, J. A., & Deutsch, D. (1963). Attention: Some Theoretical Considerations. *Psychological Review*, *70*(1), 80–90. http://doi.org/10.1037/h0039515

Dienes, Z. (2011). Bayesian versus orthodox statistics: Which side are you

on? *Perspectives on Psychological Science*. http://doi.org/10.1177/1745691611406920

Dolcos, F., Miller, B., Kragel, P., Jha, A., & McCarthy, G. (2007). Regional brain differences in the effect of distraction during the delay interval of a working memory task. *Brain Research*, *1152*, 171–181. http://doi.org/10.1016/j.brainres.2007.03.059

Dosher, B. A., & Lu, Z. L. (2000). Noise Exclusion in Spatial Attention. *Psychological Science*, *11*(2), 139–146. http://doi.org/10.1111/1467-9280.00229

Downing, P. E. (2000). Interactions between visual working memory and selective attention. *Psychological Science*, *11*(6), 467–473. http://doi.org/10.1111/1467-9280.00290

Dubé, C., Zhou, F., Kahana, M. J., & Sekuler, R. (2014). Similarity-based distortion of visual short-term memory is due to perceptual averaging. *Vision Research*, *96C*, 8–16. http://doi.org/10.1016/j.visres.2013.12.016

Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, *113*(4), 501–517.

Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., & Reitboeck, H. J. (1988). Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological Cybernetics*, *60*(2), 121–130. http://doi.org/10.1007/BF00202899

Egly, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General*, *123*(2), 161–177.

Egner, T., Monti, J. M. P., Trittschuh, E. H., Wieneke, C. A., Hirsch, J., & Mesulam, M.-M. (2008). Neural integration of top-down spatial and feature-based information in visual search. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *28*(24), 6141–6151. http://doi.org/10.1523/JNEUROSCI.1262-08.2008

Eimer, M., & Schlaghecken, F. (1998). Effects of masked stimuli on motor activation: Behavioral and electrophysiological evidence. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(6), 1737–1747. http://doi.org/10.1037/0096-1523.24.6.1737

Eimer, M., & Schlaghecken, F. (2002). Links between conscious awareness and response inhibition: Evidence from masked priming. *Psychonomic Bulletin & Review*, *9*(3), 514–520. http://doi.org/10.3758/BF03196307

Eimer, M., & Schlaghecken, F. (2003). Response facilitation and inhibition in

subliminal priming. *Biological Psychology*, *64*(1-2), 7–26. http://doi.org/10.1016/S0301-0511(03)00100-5

Eimer, M., & Velzen, J. (2002). Crossmodal links in spatial attention are mediated by supramodal control processes: Evidence from event-related potentials. *Psychophysiology*, *39*(4), 437–449. http://doi.org/10.1111/1469-8986.3940437

Fairclough, S. H., Venables, L., & Tattersall, A. (2005). The influence of task demand and learning on the psychophysiological response. *International Journal of Psychophysiology*, *56*(2), 171–184. http://doi.org/10.1016/j.ijpsycho.2004.11.003

Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, *17*(5), 738–743. http://doi.org/10.1038/nn.3689

Fisher, N. I. (1993). Statistical analysis of spherical data.

Folk, C. L., Leber, A. B., & Egeth, H. E. (2002). Made you blink! Contingent attentional capture produces a spatial blink. *Perception & Psychophysics*, *64*(5), 741–753.

Fougnie, D., Asplund, C. L., & Marois, R. (2010). What are the units of storage in visual working memory? *Journal of Vision*, *10*(12), 27–27. http://doi.org/10.1167/10.12.27

Fougnie, D., Cormiea, S. M., Kanabar, A., & Alvarez, G. A. (2016). Strategic Trade-Offs Between Quantity and Quality in Working Memory. *Journal of Experimental Psychology: Human Perception and Performance*, 1–12. http://doi.org/10.1037/xhp0000211

Fougnie, D., Suchow, J. W., & Alvarez, G. A. (2012a). Variability in the quality of visual working memory. *Nature Communications*, *3*, 1229. http://doi.org/10.1038/ncomms2237

Fougnie, D., Suchow, J. W., & Alvarez, G. A. (2012b). Variability in the quality of visual working memory. *Nature Communications*, *3*, 1229. http://doi.org/10.1038/ncomms2237

Foxe, J. J., Simpson, G. V., Ahlfors, S. P., & Saron, C. D. (2005). Biasing the brain's attentional set: I. cue driven deployments of intersensory selective attention. *Experimental Brain Research*, *166*(3-4), 370–392. http://doi.org/10.1007/s00221-005-2378-7

Fries, P. (2009). Neuronal Gamma-Band Synchronization as a Fundamental Process in Cortical Computation. *Annual Review of Neuroscience*, *32*(1), 209–224. http://doi.org/10.1146/annurev.neuro.051508.135603

Fries, P., Reynolds, J. H., Rorie, A. E., & Desimone, R. (2001a). Modulation of oscillatory neuronal synchronization by selective visual attention.

Science, *291*(5508), 1560–1563. http://doi.org/10.1126/science.291.5508.1560

Fries, P., Reynolds, J. H., Rorie, A. E., & Desimone, R. (2001b). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, *291*(5508), 1560–1563. http://doi.org/10.1126/science.291.5508.1560

Fries, P., Reynolds, J. H., Rorie, A. E., & Desimone, R. (2015). Modulation of Oscillatory Neuronal Synchronization by Selective Visual Attention. *Science*, 1–5.

Fries, P., Womelsdorf, T., Oostenveld, R., & Desimone, R. (2008). The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *28*(18), 4823–4835. http://doi.org/10.1523/JNEUROSCI.4499-07.2008

Gayet, S., Paffen, C. L. E., & Van der Stigchel, S. (2013a). Information matching the content of visual working memory is prioritized for conscious access. *Psychological Science*, *24*(12), 2472–2480. http://doi.org/10.1177/0956797613495882

Gayet, S., Van der Stigchel, S., & Paffen, C. L. E. (2013b). Seeing is believing: Utilization of subliminal symbols requires a visible relevant context. *Attention, Perception, & Psychophysics*, *76*(2), 1–19. http://doi.org/10.3758/s13414-013-0580-4

Gazzaley, A., & Nobre, A. C. (2012). Top-down modulation: bridging selective attention and working memory. *Trends in Cognitive Sciences*, *16*(2), 129–135. http://doi.org/10.1016/j.tics.2011.11.014

Giesbrecht, B., Woldorff, M. G., Song, A. W., & Mangun, G. R. (2003). Neural mechanisms of top-down control during spatial and feature attention. *NeuroImage*, *19*(3), 496–512.

Gilbert, C. D., & Wiesel, T. N. (1990). The influence of contextual stimuli on the orientation selectivity of cells in primary visual cortex of the cat. *Vision Research*, *30*(11), 1689–1701.

Girshick, A. R., Landy, M. S., & Simoncelli, E. P. (2011). Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, *14*(7), 926–932. http://doi.org/10.1038/nn.2831

Gorgoraptis, N., Catalao, R. F. G., Bays, P. M., & Husain, M. (2011). Dynamic updating of working memory resources for visual objects. *Journal of Neuroscience*, *31*(23), 8502–8511. http://doi.org/10.1523/JNEUROSCI.0208-11.2011

Gould, I. C., Rushworth, M. F., & Nobre, A. C. (2011). Indexing the graded allocation of visuospatial attention using anticipatory alpha oscillations. *Journal of Neurophysiology*, *105*(3), 1318–1326. http://doi.org/10.1152/jn.00653.2010

Green, D. M., & Swets, J. A. (1974). Signal Detection Theory and Psychophysics (Krieger, New York).

Greenberg, A. S., Esterman, M., Wilson, D., Serences, J. T., & Yantis, S. (2010). Control of spatial and feature-based attention in frontoparietal cortex. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *30*(43), 14330–14339. http://doi.org/10.1523/JNEUROSCI.4248-09.2010

Greenwald, A. G., Draine, S. C., & Abrams, R. L. (1996). Three Cognitive Markers of Unconscious Semantic Activation. *Science*, *273*(5282), 1699–1702. http://doi.org/10.1126/science.273.5282.1699

Gregoriou, G. G., Gotts, S. J., Zhou, H., & Desimone, R. (2009a). High-Frequency, Long-Range Coupling Between Prefrontal and Visual Cortex During Attention. *Science*, *324*(5931), 1207–1210. http://doi.org/10.1126/science.1171402

Gregoriou, G. G., Gotts, S. J., Zhou, H., & Desimone, R. (2009b). Long-range neural coupling through synchronization with attention. *Progress in Brain Research*, *176*, 35–45. http://doi.org/10.1016/S0079-6123(09)17603-3

Griffin, I. C., & Nobre, A. C. (2003). Orienting attention to locations in internal representations. *Journal of Cognitive Neuroscience*, *15*(8), 1176–1194. http://doi.org/10.1162/089892903322598139

Gruber, T., & Müller, M. M. (2005). Oscillatory brain activity dissociates between associative stimulus content in a repetition priming task in the human EEG. *Cerebral Cortex (New York, N.Y. : 1991)*, *15*(1), 109–116. http://doi.org/10.1093/cercor/bhh113

Gruber, T., Müller, M. M., Keil, A., & Elbert, T. (1999). Selective visual-spatial attention alters induced gamma band responses in the human EEG. *Clinical Neurophysiology*, *110*(12), 2074–2085.

Gunseli, E., van Moorselaar, D., Meeter, M., & Olivers, C. N. L. (2015). The reliability of retro-cues determines the fate of noncued visual working memory representations. *Psychonomic Bulletin & Review*, *22*(5), 1334–1341. http://doi.org/10.3758/s13423-014-0796-x

Gutnisky, D. A., Beaman, C. B., Lew, S. E., & Dragoi, V. (2016). Spontaneous Fluctuations in Visual Cortical Responses Influence Population Coding Accuracy. *Cerebral Cortex*, bhv312. http://doi.org/10.1093/cercor/bhv312

Haenny, P. E., & Schiller, P. H. (1988). State dependent activity in monkey visual cortex. I. Single cell activity in V1 and V4 on visual tasks. *Experimental Brain Research*, *69*(2), 225–244.

Haenschel, C., Baldeweg, T., Croft, R. J., Whittington, M., & Gruzelier, J. (2000). Gamma and beta frequency oscillations in response to novel auditory stimuli: A comparison of human electroencephalogram (EEG) data with in vitro models. *Proceedings of the National Academy of Sciences of the United States of America*, *97*(13), 7645–7650. http://doi.org/10.1073/pnas.120162397

Hall, S. D., Holliday, I. E., Hillebrand, A., Singh, K. D., Furlong, P. L., Hadjipapas, A., & Barnes, G. R. (2005). The missing link: analogous human and primate cortical gamma oscillations. *NeuroImage*, *26*(1), 13–17. http://doi.org/10.1016/j.neuroimage.2005.01.009

Hannula, D. E., Simons, D. J., & Cohen, N. J. (2005). Imaging implicit perception: promise and pitfalls. *Nature Reviews Neuroscience*, *6*(3), 247–255. http://doi.org/10.1038/nrn1630

Hanslmayr, S., Aslan, A., Staudigl, T., Klimesch, W., Herrmann, C. S., & Bäuml, K.-H. (2007). Prestimulus oscillations predict visual perception performance between and within subjects. *NeuroImage*, *37*(4), 1465–1473. http://doi.org/10.1016/j.neuroimage.2007.07.011

Harms, L., & Bundesen, C. (1983). Color segregation and selective attention in a nonsearch task. *Perception & Psychophysics*, *33*(1), 11–19.

Hassin, R. R., Bargh, J. A., Engell, A. D., & McCulloch, K. C. (2009). Implicit working memory. *Consciousness and Cognition*, *18*(3), 665–678. http://doi.org/10.1016/j.concog.2009.04.003

Heinze, H. J., Mangun, G. R., Burchert, W., Hinrichs, H., Scholz, M., Münte, T. F., et al. (1994). Combined spatial and temporal imaging of brain activity during visual selective attention in humans. *Nature*, *372*(6506), 543–546. http://doi.org/10.1038/372543a0

Henrie, J. A., & Shapley, R. (2005). LFP Power Spectra in V1 Cortex: The Graded Effect of Stimulus Contrast. *Journal of Neurophysiology*, *94*(1), 479–490. http://doi.org/10.1152/jn.00919.2004

Herrmann, C. S., Munk, M. H. J., & Engel, A. K. (2004). Cognitive functions of gamma-band activity: memory match and utilization. *Trends in Cognitive Sciences*, *8*(8), 347–355. http://doi.org/10.1016/j.tics.2004.06.006

Honkanen, R., Rouhinen, S., Wang, S. H., Palva, J. M., & Palva, S. (2014). Gamma Oscillations Underlie the Maintenance of Feature-Specific Information and the Contents of Visual Working Memory. *Cerebral*

Cortex, bhu263. http://doi.org/10.1093/cercor/bhu263

Hopfinger, J. B., Buonocore, M. H., & Mangun, G. R. (2000). The neural mechanisms of top-down attentional control. *Nature Neuroscience*, *3*(3), 284–291. http://doi.org/10.1038/72999

Huang, J., & Sekuler, R. (2010). Distortions in recall from visual memory: two classes of attractors at work. *Journal of Vision*, *10*(2), 24.1–27. http://doi.org/10.1167/10.2.24

Huang, L. (2010). Characterizing the nature of visual conscious access: the distinction between features and locations. *Journal of Vision*, *10*(10), 24–24. http://doi.org/10.1167/10.10.24

Huang, L., & Dobkins, K. R. (2005). Attentional effects on contrast discrimination in humans: evidence for both contrast gain and response gain. *Vision Research*, *45*(9), 1201–1212. http://doi.org/10.1016/j.visres.2004.10.024

James, W. (1890). The principles of psychology. New York: Holt.

Jehee, J. F. M., Brady, D. K., & Tong, F. (2011). Attention improves encoding of task-relevant features in the human visual cortex. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *31*(22), 8210–8219. http://doi.org/10.1523/JNEUROSCI.6153-09.2011

Jensen, O., & Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: gating by inhibition. *Frontiers in Human Neuroscience*, *4*, 186. http://doi.org/10.3389/fnhum.2010.00186

Jensen, O., Kaiser, J., & Lachaux, J.-P. (2007). Human gamma-frequency oscillations associated with attention and memory. *Trends in Neurosciences*, *30*(7), 317–324. http://doi.org/10.1016/j.tins.2007.05.001

Jha, A. P., Fabian, S. A., & Aguirre, G. K. (2004). The role of prefrontal cortex in resolving distractor interference. *Cognitive, Affective, & Behavioral Neuroscience*, *4*(4), 517–527. http://doi.org/10.3758/CABN.4.4.517

Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2009). A Dynamic Neural Field Model of Visual Working Memory and Change Detection. *Psychological Science*, *20*(5), 568–577. http://doi.org/10.1111/j.1467-9280.2009.02329.x

Kanwisher, N., & Wojciulik, E. (2000). Visual attention: insights from brain imaging. *Nature Reviews Neuroscience*, *1*(2), 91–100. http://doi.org/10.1038/35039043

Kastner, S., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1998).

Mechanisms of directed attention in the human extrastriate cortex as revealed by functional MRI. *Science*, *282*(5386), 108–111.

Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1999). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron*, *22*(4), 751–761.

Kaufmann, J. M., Schweinberger, S. R., & MikeBurton, A. (2009). N250 ERP Correlates of the Acquisition of Face Representations across Different Images. *Dx.Doi.org*, *21*(4), 625–641. http://doi.org/10.1162/jocn.2009.21080

Kelly, S. P., Lalor, E. C., Reilly, R. B., & Foxe, J. J. (2006). Increases in alpha oscillatory power reflect an active retinotopic mechanism for distracter suppression during sustained visuospatial attention. *Journal of Neurophysiology*, *95*(6), 3844–3851. http://doi.org/10.1152/jn.01234.2005

Kiefer, M., Ansorge, U., Haynes, J.-D., Hamker, F., Mattler, U., Verleger, R., & Niedeggen, M. (2011). Neuro-cognitive mechanisms of conscious and unconscious visual perception: From a plethora of phenomena to general principles. *Advances in Cognitive Psychology*, *7*(-1), 55–67. http://doi.org/10.2478/v10053-008-0090-4

Kingstone, A. (2007). Combining Expectancies. *The Quarterly Journal of Experimental Psychology*, *44*(1), 69–104. http://doi.org/10.1080/14640749208401284

Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? Perception 36 ECVP Abstract Supplement.

Klimesch, W. (2012). α-band oscillations, attention, and controlled access to stored information. *Trends in Cognitive Sciences*, *16*(12), 606–617. http://doi.org/10.1016/j.tics.2012.10.007

Klimesch, W., Sauseng, P., & Hanslmayr, S. (2007). EEG alpha oscillations: the inhibition-timing hypothesis. *Brain Research Reviews*, *53*(1), 63–88. http://doi.org/10.1016/j.brainresrev.2006.06.003

Klimesch, W., Schimke, H., & Pfurtscheller, G. (1993). Alpha frequency, cognitive load and memory performance. *Brain Topography*, *5*(3), 241–251.

Koelewijn, L., Rich, A. N., Muthukumaraswamy, S. D., & Singh, K. D. (2013). Spatial attention increases high-frequency gamma synchronisation in human medial visual cortex. *NeuroImage*, *79*(C), 295–303. http://doi.org/10.1016/j.neuroimage.2013.04.108

Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: a critical review of visual masking. *Philosophical*

Transactions of the Royal Society of London. Series B, Biological Sciences, *362*(1481), 857–875. http://doi.org/10.1098/rstb.2007.2093

Kuo, B.-C., Stokes, M. G., & Nobre, A. C. (2012). Attention modulates maintenance of representations in visual short-term memory. *Journal of Cognitive Neuroscience*, *24*(1), 51–60. http://doi.org/10.1162/jocn_a_00087

Lalonde, J., & Chaudhuri, A. (2002). Task-dependent transfer of perceptual to memory representations during delayed spatial frequency discrimination. *Vision Research*, *42*(14), 1759–1769. http://doi.org/10.1016/S0042-6989(02)00111-6

Landau, A. N., Esterman, M., Robertson, L. C., Bentin, S., & Prinzmetal, W. (2007). Different effects of voluntary and involuntary attention on EEG activity in the gamma band. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *27*(44), 11986–11990. http://doi.org/10.1523/JNEUROSCI.3092-07.2007

Landman, R., Spekreijse, H., & Lamme, V. A. F. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, *43*(2), 149–164.

Lee, D. K., Koch, C., & Braun, J. (1997). Spatial vision thresholds in the near absence of attention. *Vision Research*, *37*(17), 2409–2418.

Lee, J., & Maunsell, J. H. R. (2010). The effect of attention on neuronal responses to high and low contrast stimuli. *Journal of Neurophysiology*, *104*(2), 960–971. http://doi.org/10.1152/jn.01019.2009

Lepsien, J., & Nobre, A. C. (2006). Cognitive control of attention in the human brain: insights from orienting attention to mental representations. *Brain Research*, *1105*(1), 20–31. http://doi.org/10.1016/j.brainres.2006.03.033

Lepsien, J., Griffin, I. C., Devlin, J. T., & Nobre, A. C. (2005). Directing spatial attention in mental representations: Interactions between attentional orienting and working-memory load. *NeuroImage*, *26*(3), 733–743. http://doi.org/10.1016/j.neuroimage.2005.02.026

Levitt, J. B., & Lund, J. S. (1997). Contrast dependence of contextual effects in primate visual cortex. *Nature*, *387*(6628), 73–76. http://doi.org/10.1038/387073a0

Liu, T., Slotnick, S. D., Serences, J. T., & Yantis, S. (2003). Cortical mechanisms of feature-based attentional control. *Cerebral Cortex (New York, N.Y. : 1991)*, *13*(12), 1334–1343. http://doi.org/10.1093/cercor/bhg080

Lu, Z. L., & Dosher, B. A. (1998). External noise distinguishes attention mechanisms. *Vision Research*, *38*(9), 1183–1198.

Lu, Z. L., & Dosher, B. A. (2008). Characterizing observers using external noise and observer models: assessing internal representations with external noise. *Psychological Review*. http://doi.org/10.1037/0033-295X.115.1.44

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

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. http://doi.org/10.1038/363345a0

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

Mack, A., & Rock, I. (1998). Inattentional blindness. http://doi.org/10.1016/pic.2015.12.17.2502

Macmillan, N. A., & Creelman, C. D. (1991). *Detection Theory: A User's Guide*. Cambridge UP, Cambridge.

Macmillan, N. A., & Creelman, C. D. (2005). *Detection Theory: A User's Guide*. Lawrence Erlbaum Associates. New York.

Magnussen, S. (2000). Low-level memory processes in vision. *Trends in Neurosciences*, *23*(6), 247–251.

Magnussen, S., Greenlee, M. W., Asplund, R., & Dyrnes, S. (1991). Stimulus-specific mechanisms of visual short-term memory. *Vision Research*, *31*(7-8), 1213–1219.

Makovski, T. (2012). Are multiple visual short-term memory storages necessary to explain the retro-cue effect? *Psychonomic Bulletin & Review*, *19*(3), 470–476. http://doi.org/10.3758/s13423-012-0235-9

Makovski, T., & Jiang, Y. V. (2007). Distributing versus focusing attention in visual short-term memory. *Psychonomic Bulletin & Review*, *14*(6), 1072–1078. http://doi.org/10.3758/BF03193093

Makovski, T., & Pertzov, Y. (2015). Attention and memory protection: Interactions between retrospective attention cueing and interference. *The Quarterly Journal of Experimental Psychology*, *68*(9), 1735–1743. http://doi.org/10.1080/17470218.2015.1049623

Makovski, T., Sussman, R., & Jiang, Y. V. (2008). Orienting attention in visual working memory reduces interference from memory probes.

Journal of Experimental Psychology: Learning, Memory, and Cognition, *34*(2), 369–380. http://doi.org/10.1037/0278-7393.34.2.369

Martinez, A., Teder-Sälejärvi, W., Vazquez, M., Molholm, S., Foxe, J. J., Javitt, D. C., et al. (2006). Objects are highlighted by spatial attention. *Journal of Cognitive Neuroscience*, *18*(2), 298–310. http://doi.org/10.1162/089892906775783642

Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curbio*, *14*(9), 744–751. http://doi.org/10.1016/j.cub.2004.04.028

Martínez-Trujillo, J., & Treue, S. (2002). Attentional modulation strength in cortical area MT depends on stimulus contrast. *Neuron*, *35*(2), 365–370.

Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., & Ro, T. (2009). To see or not to see: prestimulus alpha phase predicts visual awareness. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *29*(9), 2725–2732. http://doi.org/10.1523/JNEUROSCI.3963-08.2009

Matsukura, M., Luck, S. J., & Vecera, S. P. (2007). Attention effects during visual short-term memory maintenance: protection or prioritization? *Perception & Psychophysics*, *69*(8), 1422–1434.

Matthey, L., Bays, P. M., & Dayan, P. (2015). A probabilistic palimpsest model of visual short-term memory. *PLoS Computational Biology*, *11*(1), e1004003. http://doi.org/10.1371/journal.pcbi.1004003

Maunsell, J. H. R., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, *29*(6), 317–322. http://doi.org/10.1016/j.tins.2006.04.001

McAdams, C. J., & Maunsell, J. H. (1999). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *Journal of Neuroscience*, *19*(1), 431–441.

McAdams, C. J., & Maunsell, J. H. (2000). Attention to both space and feature modulates neuronal responses in macaque area V4. *Journal of Neurophysiology*, *83*(3), 1751–1755.

Mesulam, M. M. (1990). Large-scale neurocognitive networks and distributed processing for attention, language, and memory. *Annals of Neurology*, *28*(5), 597–613. http://doi.org/10.1002/ana.410280502

Motter, B. C. (1994). Neural correlates of attentive selection for color or luminance in extrastriate area V4. *Journal of Neuroscience*, *14*(4), 2178–2189.

Murray, A. M., Nobre, A. C., Astle, D. E., & Stokes, M. G. (2012). Lacking

control over the trade-off between quality and quantity in visual short-term memory. *PLoS ONE*, *7*(8), e41223. http://doi.org/10.1371/journal.pone.0041223

Murray, A. M., Nobre, A. C., Clark, I. A., Cravo, A. M., & Stokes, M. G. (2013). Attention restores discrete items to visual short-term memory. *Psychological Science*, *24*(4), 550–556. http://doi.org/10.1177/0956797612457782

Müller, M. M., & Keil, A. (2004). Neuronal synchronization and selective color processing in the human brain. *Journal of Cognitive Neuroscience*, *16*(3), 503–522. http://doi.org/10.1162/089892904322926827

Müller, M. M., Teder-Sälejärvi, W., & Hillyard, S. A. (1998). The time course of cortical facilitation during cued shifts of spatial attention. *Nature Neuroscience*, *1*(7), 631–634. http://doi.org/10.1038/2865

Myers, N. E., Stokes, M. G., Walther, L., & Nobre, A. C. (2014). Oscillatory brain state predicts variability in working memory. *Journal of Neuroscience*, *34*(23), 7735–7743. http://doi.org/10.1523/JNEUROSCI.4741-13.2014

Myers, N. E., Walther, L., Wallis, G., Stokes, M. G., & Nobre, A. C. (2015). Temporal Dynamics of Attention during Encoding versus Maintenance of Working Memory: Complementary Views from Event-related Potentials and Alpha-band Oscillations. *Journal of Cognitive Neuroscience*, *27*(3), 492–508. http://doi.org/10.1162/jocn_a_00727

Neisser, U. (1967). *Cognitive psychology.* East Norwalk, CT, US: Appleton-Century-Crofts.

Nemes, V. A., Parry, N. R. A., Whitaker, D., & McKeefry, D. J. (2012). The retention and disruption of color information in human short-term visual memory. *Journal of Vision*, *12*(1), 26–26. http://doi.org/10.1167/12.1.26

Neumann, O., & Klotz, W. (1994). Motor responses to nonreportable, masked stimuli: Where is the limit of direct parameter specification. *Attention and Performance XV: Conscious and nonconscious information processing,* e.d. C. Umiltà & M. Moscovitch. MIT Press.

Nicol, R. M., Chapman, S. C., Vértes, P. E., Nathan, P. J., Smith, M. L., Shtyrov, Y., & Bullmore, E. T. (2012). Fast reconfiguration of high-frequency brain networks in response to surprising changes in auditory input. *Journal of Neurophysiology*, *107*(5), 1421–1430. http://doi.org/10.1152/jn.00817.2011

Nobre, A. C. (2004). Probing the flexibility of attentional orienting in the human brain. (M. I. Posner, Ed.) (pp. 157–179). New York: Guilford Press. http://doi.org/10.1162/jocn.2006.18.5.715

Nobre, A. C., Griffin, I. C., & Rao, A. (2007). Spatial attention can bias

search in visual short-term memory. *Frontiers in Human Neuroscience*, *1*, 4. http://doi.org/10.3389/neuro.09.004.2007

Nobre, A. C., Sebestyen, G. N., Gitelman, D. R., Mesulam, M. M., Frackowiak, R. S., & Frith, C. D. (1997). Functional localization of the system for visuospatial attention using positron emission tomography. *Brain*, *120 (Pt 3)*, 515–533.

Ogawa, T., & Komatsu, H. (2004). Target selection in area V4 during a multidimensional visual search task. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *24*(28), 6371–6382. http://doi.org/10.1523/JNEUROSCI.0569-04.2004

Parkes, L., Lund, J., Angelucci, A., Solomon, J. A., & Morgan, M. (2001). Compulsory averaging of crowded orientation signals in human vision. *Nature Neuroscience*, *4*(7), 739–744. http://doi.org/10.1038/89532

Pasternak, T., & Zaksas, D. (2003). Stimulus specificity and temporal dynamics of working memory for visual motion. *Journal of Neurophysiology*, *90*(4), 2757–2762. http://doi.org/10.1152/jn.00422.2003

Pertzov, Y., Bays, P. M., Joseph, S., & Husain, M. (2013). Rapid forgetting prevented by retrospective attention cues. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(5), 1224–1231. http://doi.org/10.1037/a0030947

Pesaran, B., Pezaris, J. S., Sahani, M., Mitra, P. P., & Andersen, R. A. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nature Neuroscience*, *5*(8), 805–811. http://doi.org/10.1038/nn890

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

Posner, M. I. (1994). Attention: the mechanisms of consciousness. *Proceedings of the National Academy of Sciences of the United States of America*, *91*(16), 7398–7403.

Posner, M. I., Cohen, Y., & Rafal, R. D. (1982). Neural systems control of spatial orienting. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *298*(1089), 187–198.

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.

Posner, M. I., Walker, J. A., Friedrich, F. J., & Rafal, R. D. (1984). Effects of parietal injury on covert orienting of attention. *Journal of Neuroscience*,

4(7), 1863–1874.

Postle, B. R. (2005). Delay-period Activity in the Prefrontal Cortex: One Function Is Sensory Gating. *Journal of Cognitive Neuroscience*, *17*(11), 1679–1690. http://doi.org/10.1162/089892905774589208

Pratte, M. S., & Rouder, J. N. (2009). A task-difficulty artifact in subliminal priming. *Attention, Perception, & Psychophysics*, *71*(6), 1276–1283. http://doi.org/10.3758/APP.71.6.1276

Pulvermüller, F., Keil, A., & Elbert, T. (1999). High-frequency brain activity: perception or active memory? *Trends in Cognitive Sciences*, *3*(7), 250–252.

Rademaker, R. L., Bloem, I. M., De Weerd, P., & Sack, A. T. (2015). The impact of interference on short-term memory for visual orientation. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(6), 1650–1665. http://doi.org/10.1037/xhp0000110

Ramsøy, T. Z., & Overgaard, M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences*, *3*(1), 1–23. http://doi.org/10.1023/B:PHEN.0000041900.30172.e8

Ramsøy, T. Z., Rams, & Overgaard, M. (2004). Introspection and subliminal perception. *Phenomenology and the Cognitive Sciences*, *3*, 1–23.

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. http://doi.org/10.1037/0096-1523.18.3.849

Reingold, E. M. (2004). Unconscious perception: assumptions and interpretive difficulties. *Consciousness and Cognition*, *13*(1), 117–122. http://doi.org/10.1016/j.concog.2003.11.002

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To See or not to See: The Need for Attention to Perceive Changes in Scenes. *Psychological Science*, *8*(5), 368–373. http://doi.org/10.1111/j.1467-9280.1997.tb00427.x

Rerko, L., & Oberauer, K. (2013). Focused, unfocused, and defocused information in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(4), 1075–1096. http://doi.org/10.1037/a0031172

Rerko, L., Souza, A. S., & Oberauer, K. (2014). Retro-cue benefits in working memory without sustained focal attention. *Memory & Cognition*, *42*(5), 712–728. http://doi.org/10.3758/s13421-013-0392-8

Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of

attention. *Neuron*, *61*(2), 168–185. http://doi.org/10.1016/j.neuron.2009.01.002

Reynolds, J. H., Pasternak, T., & Desimone, R. (2000). Attention increases sensitivity of V4 neurons. *Neuron*, *26*(3), 703–714.

Rock, I., & Gutman, D. (1981). The effect of inattention on form perception. *Journal of Experimental Psychology: Human Perception and Performance, 7(2), 275 - 285.*

Saenz, M., Buracas, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience*, *5*(7), 631–632. http://doi.org/10.1038/nn876

Saenz, M., Buracas, G. T., & Boynton, G. M. (2003). Global feature-based attention for motion and color. *Vision Research*, *43*(6), 629–637. http://doi.org/10.1016/S0042-6989(02)00595-3

Sauseng, P., Klimesch, W., Stadler, W., Schabus, M., Doppelmayr, M., Hanslmayr, S., et al. (2005). A shift of visual spatial attention is selectively associated with human EEG alpha activity. *European Journal of Neuroscience*, *22*(11), 2917–2926. http://doi.org/10.1111/j.1460-9568.2005.04482.x

Schenkluhn, B., Ruff, C. C., Heinen, K., & Chambers, C. D. (2008). Parietal stimulation decouples spatial and feature-based attention. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *28*(44), 11106–11110. http://doi.org/10.1523/JNEUROSCI.3591-08.2008

Schweinberger, S. R., Pickering, E. C., Jentzsch, I., Burton, A. M., & Kaufmann, J. M. (2002). Event-related brain potential evidence for a response of inferior temporal cortex to familiar face repetitions. *Cognitive Brain Research*, *14*(3), 398–409. http://doi.org/10.1016/S0926-6410(02)00142-8

Serences, J. T., & Boynton, G. M. (2007). Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron*, *55*(2), 301–312. http://doi.org/10.1016/j.neuron.2007.06.015

Serences, J. T., Saproo, S., Scolari, M., Ho, T., & Muftuler, L. T. (2009). Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. *NeuroImage*, *44*(1), 223–231. http://doi.org/10.1016/j.neuroimage.2008.07.043

Serences, J. T., Shomstein, S., Leber, A. B., Golay, X., Egeth, H. E., & Yantis, S. (2005). Coordination of voluntary and stimulus-driven attentional control in human cortex. *Psychological Science*, *16*(2), 114–122. http://doi.org/10.1111/j.0956-7976.2005.00791.x

Sergent, C., Wyart, V., Babo-Rebelo, M., Cohen, L., Naccache, L., & Tallon-Baudry, C. (2013). Cueing Attention after the Stimulus Is Gone Can

Retrospectively Trigger Conscious Perception. *Current Biology*, *23*(2), 150–155. http://doi.org/10.1016/j.cub.2012.11.047

Shulman, G. L., d'Avossa, G., Tansy, A. P., & Corbetta, M. (2002). Two attentional processes in the parietal lobe. *Cerebral Cortex (New York, N.Y. : 1991)*, *12*(11), 1124–1131.

Siegel, M., Donner, T. H., Oostenveld, R., Fries, P., & Engel, A. K. (2008). Neuronal synchronization along the dorsal visual pathway reflects the focus of spatial attention. *Neuron*, *60*(4), 709–719. http://doi.org/10.1016/j.neuron.2008.09.010

Silvanto, J., & Soto, D. (2012). Causal evidence for subliminal percept-to-memory interference in early visual cortex. *NeuroImage*. http://doi.org/10.1016/j.neuroimage.2011.07.062

Singer, W., & Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, *18*(1), 555–586. http://doi.org/10.1146/annurev.ne.18.030195.003011

Sligte, I. G., Scholte, H. S., & Lamme, V. A. F. (2008). Are There Multiple Visual Short-Term Memory Stores? *PLoS ONE*, *3*(2), e1699. http://doi.org/10.1371/journal.pone.0001699

Slotnick, S. D., Schwarzbach, J., & Yantis, S. (2003). Attentional inhibition of visual processing in human striate and extrastriate cortex. *NeuroImage*, *19*(4), 1602–1611. http://doi.org/10.1016/S1053-8119(03)00187-3

Smith, A. T., Singh, K. D., & Greenlee, M. W. (2000). Attentional suppression of activity in the human visual cortex. *Neuroreport*, *11*(2), 271–277.

Sneve, M. H., Sreenivasan, K. K., Alnæs, D., Endestad, T., & Magnussen, S. (2015). Short-term retention of visual information: Evidence in support of feature-based attention as an underlying mechanism. *Neuropsychologia*, *66*, 1–9. http://doi.org/10.1016/j.neuropsychologia.2014.11.004

Snyder, A. C., & Foxe, J. J. (2010). Anticipatory attentional suppression of visual features indexed by oscillatory alpha-band power increases: a high-density electrical mapping study. *Journal of Neuroscience*, *30*(11), 4024–4032. http://doi.org/10.1523/JNEUROSCI.5684-09.2010

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. http://doi.org/10.1016/j.tics.2008.05.007

Soto, D., Mäntylä, T., & Silvanto, J. (2011). Working memory without consciousness. *Current Biology*. http://doi.org/10.1016/j.cub.2011.09.049

Souza, A. S., Rerko, L., & Oberauer, K. (2016). Getting More From Visual Working Memory: Retro-Cues Enhance Retrieval and Protect From Visual Interference. *Journal of Experimental Psychology: Human Perception and Performance*, 1–21. http://doi.org/10.1037/xhp0000192

Souza, A. S., Rerko, L., Lin, H.-Y., & Oberauer, K. (2014). Focused attention improves working memory: implications for flexible-resource and discrete-capacity models. *Attention, Perception, & Psychophysics*, *76*(7), 2080–2102. http://doi.org/10.3758/s13414-014-0687-2

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, *74*(11), 1–29. http://doi.org/10.1037/h0093759

Sreenivasan, K. K., & Jha, A. P. (2007). Selective attention supports working memory maintenance by modulating perceptual processing of distractors. *Journal of Cognitive Neuroscience*, *19*(1), 32–41. http://doi.org/10.1162/jocn.2007.19.1.32

Sreenivasan, K. K., Sambhara, D., & Jha, A. P. (2011). Working memory templates are maintained as feature-specific perceptual codes. *Journal of Neurophysiology*, *106*(1), 115–121. http://doi.org/10.1152/jn.00776.2010

Stein, T., Kaiser, D., & Hesselmann, G. (2016). Can working memory be non-conscious? *Neuroscience of Consciousness*, *2016*(1), niv011–3. http://doi.org/10.1093/nc/niv011

Steinmetz, P. N., Roy, A., Fitzgerald, P. J., Hsiao, S. S., Johnson, K. O., & Niebur, E. (2000). Attention modulates synchronized neuronal firing in primate somatosensory cortex. *Nature*, *404*(6774), 187–190. http://doi.org/10.1038/35004588

Tallon, C., Bertrand, O., Bouchet, P., & Pernier, J. (1995). Gamma-range activity evoked by coherent visual stimuli in humans. *European Journal of Neuroscience*, *7*(6), 1285–1291.

Tallon-Baudry, C. (2009). The roles of gamma-band oscillatory synchrony in human visual cognition. *Front Biosci*., *14, 321-332.*

Tallon-Baudry, C., Bertrand, O., Delpuech, C., & Permier, J. (1997). Oscillatory gamma-band (30-70 Hz) activity induced by a visual search task in humans. *Journal of Neuroscience*, *17*(2), 722–734.

Tallon-Baudry, C., Bertrand, O., Delpuech, C., & Pernier, J. (1996). Stimulus specificity of phase-locked and non-phase-locked 40 Hz visual responses in human. *Journal of Neuroscience*, *16*(13), 4240–4249.

Tallon-Baudry, C., Bertrand, O., Peronnet, F., & Pernier, J. (1998). Induced gamma-band activity during the delay of a visual short-term memory

task in humans. *Journal of Neuroscience*, *18*(11), 4244–4254.

Tallon-Baudry, C., Kreiter, A., & Bertrand, O. (1999). Sustained and transient oscillatory responses in the gamma and beta bands in a visual short-term memory task in humans. *Visual Neuroscience*, *16*(03), 449–459.

Tanaka, J. W., Curran, T., Porterfield, A. L., & Collins, D. (2006). Activation of Preexisting and Acquired Face Representations: The N250 Event-related Potential as an Index of Face Familiarity. *Dx.Doi.org*, *18*(9), 1488–1497. http://doi.org/10.1162/jocn.2006.18.9.1488

Taulu, S., Simola, J., & Kajola, M. (2005). Applications of the Signal Space Separation Method. *Signal Processing, IEEE Transactions on*, *53*(9), 3359–3372. http://doi.org/10.1109/TSP.2005.853302

Thibault, L., van den Berg, R., Cavanagh, P., & Sergent, C. (2016). Retrospective Attention Gates Discrete Conscious Access to Past Sensory Stimuli. *PLoS ONE*, *11*(2), e0148504. http://doi.org/10.1371/journal.pone.0148504

Thut, G., Nietzel, A., Brandt, S. A., & Pascual-Leone, A. (2006). Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *26*(37), 9494–9502. http://doi.org/10.1523/JNEUROSCI.0875-06.2006

Todorovic, A., van Ede, F., Maris, E., & de Lange, F. P. (2011). Prior expectation mediates neural adaptation to repeated sounds in the auditory cortex: an MEG study. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *31*(25), 9118–9123. http://doi.org/10.1523/JNEUROSCI.1425-11.2011

Tootell, R. B., Hadjikhani, N., Hall, E. K., Marrett, S., Vanduffel, W., Vaughan, J. T., & Dale, A. M. (1998). The retinotopy of visual spatial attention. *Neuron*, *21*(6), 1409–1422.

Treue, S., & Martínez Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, *399*(6736), 575–579. http://doi.org/10.1038/21176

Tulving, E., & Schachter, D. L. (2013). Priming and Human Memory Systems. *Science*, *247*(4940), 301.

van den Berg, R., Awh, E., & Ma, W. J. (2014). Factorial comparison of working memory models. *Psychological Review*, *121*(1), 124–149. http://doi.org/10.1037/a0035234

van den Berg, R., Shin, H., Chou, W.-C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory

limitations. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(22), 8780–8785. http://doi.org/10.1073/pnas.1117465109

van Dijk, H., Schoffelen, J.-M., Oostenveld, R., & Jensen, O. (2008). Prestimulus oscillatory activity in the alpha band predicts visual discrimination ability. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *28*(8), 1816–1823. http://doi.org/10.1523/JNEUROSCI.1853-07.2008

van Moorselaar, D., Gunseli, E., Theeuwes, J., & N L Olivers, C. (2015). The time course of protecting a visual memory representation from perceptual interference. *Frontiers in Human Neuroscience*, *8*, 1053. http://doi.org/10.3389/fnhum.2014.01053

Van Voorhis, S., & Hillyard, S. A. (1977). Visual evoked potentials and selective attention to points in space. *Perception & Psychophysics*, *22*(1), 54–62. http://doi.org/10.3758/BF03206080

Vanduffel, W., Tootell, R. B., & Orban, G. A. (2000). Attention-dependent suppression of metabolic activity in the early stages of the macaque visual system. *Cerebral Cortex (New York, N.Y. : 1991)*, *10*(2), 109–126.

Vecera, S. P., & Farah, M. J. (1994). Does visual attention select objects or locations? *Journal of Experimental Psychology: General*, *123*(2), 146–160. http://doi.org/10.1037/0096-3445.123.2.146

Vidal, J. R., Chaumon, M., O'Regan, J. K., & Tallon-Baudry, C. (2006). Visual grouping and the focusing of attention induce gamma-band oscillations at different frequencies in human magnetoencephalogram signals. *Journal of Cognitive Neuroscience*, *18*(11), 1850–1862. http://doi.org/10.1162/jocn.2006.18.11.1850

Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, *428*(6984), 748–751. http://doi.org/10.1038/nature02447

Wallis, G., Stokes, M., Cousijn, H., Woolrich, M., & Nobre, A. C. (2015). Frontoparietal and Cingulo-opercular Networks Play Dissociable Roles in Control of Working Memory. *Journal of Cognitive Neuroscience*, *27*(10), 2019–2034. http://doi.org/10.1162/jocn_a_00838

White, A. L., Rolfs, M., & Carrasco, M. (2015). Stimulus competition mediates the joint effects of spatial and feature-based attention. *Journal of Vision*, *15*(14), 7–7. http://doi.org/10.1167/15.14.7

Williams, M., Hong, S. W., Kang, M.-S., Carlisle, N. B., & Woodman, G. F. (2013). The benefit of forgetting. *Psychonomic Bulletin & Review*, *20*(2), 348–355. http://doi.org/10.3758/s13423-012-0354-3

Williford, T., & Maunsell, J. H. R. (2006). Effects of spatial attention on contrast response functions in macaque area V4. *Journal of Neurophysiology*, *96*(1), 40–54. http://doi.org/10.1152/jn.01207.2005

Woldorff, M. G., Fox, P. T., Matzke, M., Lancaster, J. L., Veeraswamy, S., Zamarripa, F., et al. (1997). Retinotopic organization of early visual spatial attention effects as revealed by PET and ERPs. *Human Brain Mapping*, *5*(4), 280–286. http://doi.org/10.1002/(SICI)1097-0193(1997)5:4<280::AID-HBM13>3.0.CO;2-I

Womelsdorf, T., & Fries, P. (2007). The role of neuronal synchronization in selective attention. *Current Opinion in Neurobiology*, *17*(2), 154–160. http://doi.org/10.1016/j.conb.2007.02.002

Womelsdorf, T., Fries, P., Mitra, P. P., & Desimone, R. (2006). Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature*, *439*(7077), 733–736. http://doi.org/10.1038/nature04258

Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *20*(6), RC63.

Wyart, V., & Tallon-Baudry, C. (2008a). Neural Dissociation between Visual Awareness and Spatial Attention. *Journal of Neuroscience*, *28*(10), 2667–2679. http://doi.org/10.1523/JNEUROSCI.4748-07.2008

Wyart, V., & Tallon-Baudry, C. (2008b). Neural Dissociation between Visual Awareness and Spatial Attention. *Journal of Neuroscience*, *28*(10), 2667–2679. http://doi.org/10.1523/JNEUROSCI.4748-07.2008

Wyart, V., Nobre, A. C., & Summerfield, C. (2012). Dissociable prior influences of signal probability and relevance on visual contrast sensitivity. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(9), 3593–3598. http://doi.org/10.1073/pnas.1120118109

Yantis, S., Schwarzbach, J., Serences, J. T., Carlson, R. L., Steinmetz, M. A., Pekar, J. J., & Courtney, S. M. (2002). Transient neural activity in human parietal cortex during spatial attention shifts. *Nature Neuroscience*, *5*(10), 995–1002. http://doi.org/10.1038/nn921

Yeshurun, Y., & Carrasco, M. (1998). Attention improves or impairs visual performance by enhancing spatial resolution. *Nature*, *396*(6706), 72–75. http://doi.org/10.1038/23936

Yoon, J. H., Curtis, C. E., & D'Esposito, M. (2006). Differential effects of distraction during working memory on delay-period activity in the

prefrontal cortex and the visual association cortex. *NeuroImage*, *29*(4), 1117–1126. http://doi.org/10.1016/j.neuroimage.2005.08.024

Yuval-Greenberg, S., Tomer, O., Keren, A. S., Nelken, I., & Deouell, L. Y. (2008). Transient Induced Gamma-Band Response in EEG as a Manifestation of Miniature Saccades. *Neuron*, *58*(3), 429–441. http://doi.org/10.1016/j.neuron.2008.03.027

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*(7192), 233–235. http://doi.org/10.1038/nature06860

Zhang, W., & Luck, S. J. (2011). The number and quality of representations in working memory. *Psychological Science*, *22*(11), 1434–1441. http://doi.org/10.1177/0956797611417006

Zoccolan, D., Cox, D. D., & DiCarlo, J. J. (2005). Multiple object response normalization in monkey inferotemporal cortex. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *25*(36), 8150–8164. http://doi.org/10.1523/JNEUROSCI.2058-05.2005

Zokaei, N., Manohar, S., Husain, M., & Feredoes, E. (2014). Causal evidence for a privileged working memory state in early visual cortex. *Journal of Neuroscience*, *34*(1), 158–162. http://doi.org/10.1523/JNEUROSCI.2899-13.2014