The role of perceptual load in orientation perception, visual masking and contextual integration

Moritz J.K. Stolte

Institute of Cognitive Neuroscience
University College London

Submitted for the degree of PhD, September 2013

I, Moritz Stolte, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Previous research of the effects of perceptual load on visual processing demonstrates reduced perceptual sensitivity and reduced neural activity for task-irrelevant stimuli under high compared to low perceptual load. However, the precise underlying mechanisms for reduced processing remain unclear. The present thesis approaches this question by assessing the interactions between low level visual stimuli and their visual context under different levels of perceptual load. The results from Chapter 2 demonstrate that perceptual load modulates visual perception of orientation not only by reducing overall signal gain but also by broadening the orientation tuning profile. These findings suggest that perceptual load not only reduces the signal strength but also reduces the extent to which the signal is discriminated from noise; thus altering the strength of contextual interaction. Chapter 3 further demonstrates the role of perceptual load in contextual interactions by establishing the effects of load on the tilt-illusion, which is thought to rely purely on interactions among orientations. The results show that high perceptual load increases the impact of context (leading to greater tilt-illusion) for subthreshold context presentation, which precludes top-down suppression of the task-irrelevant context. Chapters 4 to 5 further establish reduced efficiency of separating signal from noise with higher perceptual load, found when signal and noise are presented successively (as in backward masking). The time course of perceptual load effects shows a distinct pattern of both, more effective and longer lasting masking under high compared to low load. This pattern evolves rapidly for *pattern masking* which indicates low-level integration and demonstrates an early locus for the effects of load. It also persists at later periods for *metacontrast* and *object substitution masking* where the stimulus and mask do not spatially overlap, indicating

load effects at later processing stages. Together, the results provide mechanistic explanations for reduced perception under high perceptual load.

Table of Contents

List of Tables and Figures

Acknowledgements

First and foremost I would like to thank my PhD supervisor, Nilli Lavie, for her invaluable guidance, strong moral support and countless inspiring discussions. Thank you for being a real hands-on supervisor and devoting so much of your time to the work of PhD students in the lab. I would like to express special thanks to Bahador Bahrami for guidance with psychophysical methods. Also, a big thank you to everyone in the Lavie attention lab, past and present. Thank you Alessandro, Ana, Anna, Dana, David, Rashmi, Luke, Kate, Jake, Jean-Remi, Maha, Sophie and Todd. Thanks also to Gwenisha for helping with data collection. Lastly, I would like to thank the most important person in my life, Hyunji. Thanks for your calming support during stressful periods, without you this would have not been possible.

This research has been supported by a demonstratorship from the University College London Division of Psychology and Language Sciences.

Chapter 1

General Introduction

1.1 Preface

When attention is focused on a demanding task, ordinarily clearly visible objects can go completely unnoticed. As magicians are well aware, this occurs even when we are particularly keen to perceive as much detail as possible of the scene enfolding in front of us. This creates a false sense of omniscient perception, making the magic trick seem even more implausible and surprising. For visual perception, the magnitude of this effect (i.e. the strength of perceptual suppression of unattended items) depends on the level of perceptual load in the attended task (e.g. Cartwright-Finch & Lavie, 2007; Lavie, 2006; Macdonald & Lavie, 2008; see Lavie, 2005; 2010 for reviews). Various behavioural and neuroimaging studies have validated the effects of perceptual load on task irrelevant as well as task relevant items (e.g. in dual task paradigms), demonstrating reduced perceptual sensitivity and reduced neural responses under high load. The primary aim of the current thesis is to investigate the underlying mechanisms of this reduced perception under high perceptual load.

In real world conditions, ignored objects do not occur in isolation but are embedded within a dynamic visual context. In addition to reducing responses directly representing the unattended objects, perceptual load may also result in failure to successfully separate these objects from their context. Therefore, the observed reductions in neural activity for unattended stimuli may reflect increased noise in the stimulus representation instead of (or in addition to) a reduction of the signal itself. To quantify these mechanisms, a methodological approach was chosen that can distinguish between effects on the signal strength of perceptual representations of unattended stimuli, and effects influencing the precision of such representations. By presenting targets to be detected or discriminated in the presence of visual context and systematically varying spatial, featural and temporal relationships between target and context, the investigation of the effects of perceptual load is expanded to contextual interaction in the spatial and temporal domain.

The research in this thesis, therefore, addresses the effects of perceptual load on contextual interactions in visual perception. First, in Chapter 2 the effects of load on orientation gain and tuning are examined by varying the orientation content of a noise mask presented simultaneously with an oriented target stimulus. This reveals the precision and orientation signal strength of the underlying population response. A reduction in the ability of the visual system to separate signal and noise under conditions of high load may also indicate greater modulation of the context and its interaction with target stimuli. This is investigated in Chapter 3 by assessing the effect of perceptual load on the tilt-illusion where the context determines the perceived tilt of a target (tilt-repulsion). Here, modulation of the magnitude of the illusion may reveal the influence of perceptual load on spatial contextual integration that relies purely on orientation interactions in early visual cortex, providing further support of load effects beyond processes relating to target stimulus perception per se. In addition, considering the dynamic nature of contextual interaction, perceptual load may affect separation of signal and noise even when they occur at different points in time. In chapters 4 and 5, therefore, the effect of load on the time course of contextual interaction is examined by presenting a mask at various stimulus-onset-asynchronies (SOAs) after target presentation, while varying the spatial proximity, as well as, the shared feature content between target and mask.

The following sections will briefly review some of the relevant empirical findings that provided the starting point for the experiments in this thesis. First, the principles of perceptual load theory will be outlined, followed by a summary of the existing behavioural and neurophysiological evidence in support of the theory. Next the review focuses on neurophysiological and psychophysical research that examined the distinction between attentional facilitation (by enhancing target stimulus representations) and attentional selection (by reducing the influence of non-target stimuli). Finally, attentional effects that extend beyond processing of target information (and the *classical* receptive field) are described by examining psychophysical and neurophysiological findings and theoretical accounts for contextual and temporal integration.

1.2 Perceptual Load Theory

The amount of visual information in the world around us far exceeds the capacity of the visual system. For example, a simple calculation based on the energy consumption of human cortex suggests that at any given time only 1% of the brain's neurons can be significantly active (Lennie, 2003). For the visual system this implies that a selective process is required to occur within the visual processing stream, enabling us to perceive a subset of external stimuli. The question of how (and where) such selectivity comes about in the brain has been a focal point in the cognitive sciences for over 50 years (e.g. Broadbent, 1958). Early theoretical developments provided the starting point for a long standing dichotomy in the field: *early* vs. *late selection*. Does attention exclude the processing of irrelevant stimuli from the early levels of perceptual processing (*early selection*) or do we initially perceive all stimuli and discard irrelevant or distracting information at later stages (*late selection*)? For many years, both views have received plenty of supporting evidence, maintaining an active debate over the issue until today (see Driver, 2001 for review).

A useful approach to resolving this debate has been provided by Lavie (Lavie, 1995; Lavie, 2005) under the framework of Load theory. Lavie's Load theory suggests an important role for load in the processing of *task relevant* stimuli on the extent of *task irrelevant* distractor processing. This suggests that distractor processing is reduced at early stages when perceptual load in the task is high (*early selection*) but may be excluded at later processing stages for memory or decision making when perceptual load in the task is low (*late selection*). According to the theory, a task involving a high perceptual load (e.g., visual search for a target among many or very similar distractors) will exhaust all available resources, leaving none to "spill over" to the processing of distractors. Tasks involving low perceptual load (e.g., visual search for a target among few or very different distractors), however, may only require a small proportion of the available resources, leaving the rest to process the distractors. Importantly, any spare resources that are not devoted to the processing of task relevant stimuli are thought to automatically "spill over" to all other stimuli in the scene. Thus, according to this view, selectively processing only task relevant stimuli really depends on the perceptual load of the task and cannot be voluntarily controlled.

An important aspect of Load theory is that it distinguishes between perceptual load and general task difficulty. Perceptual load is high when the number of different distractor items is increased or, using the same number of items, target identification is made attentionally more demanding (e.g. feature-conjunction search vs. single-feature search). It has been shown previously, that perceptual resources cannot compensate for data limits imposed on perceptual identification (Norman & Bobrow, 1975). Thus, although a task becomes more demanding when sensory input is degraded (e.g. reduced target contrast or size), distractor processing is unaffected (Lavie & De Fockert, 2003). In addition, more general effects of increased task difficulty under high load have been ruled out as a possible explanation for reduced distractor processing by demonstrating that increasing the overall cognitive load of the task (e.g. loading working memory) *increases* distractor processing (Lavie, Hirst, De Fockert & Viding, 2004).

This latter finding highlights the contrasting effects of load in perception and load in working memory: While perceptual load exhausts capacity for processing relevant stimuli and thus, reduces irrelevant distractor interference, working memory load exhausts capacity for cognitive control mechanisms (for example, to reject irrelevant stimuli under low perceptual load) and thus, increases distractor interference. While these later revisions of Load theory discuss the role of working memory in detail, the current thesis will only focus on the effects of perceptual load.

1.2.1 Behavioural and neurophysiological evidence for the effects of perceptual load

Evidence for the effects of perceptual load on distractor processing has been obtained in several behavioural and neuroimaging experiments. Studies using the response competition paradigm (Eriksen & Eriksen, 1975) have typically shown that when perceptual load is low, reaction times depend on the type of distractor (congruent vs. incongruent to the target) while in situations of high perceptual load these effects are almost completely eliminated, indicating much reduced distractor processing (Figure 1.1; Lavie, 1995; 2000; Lavie & Cox, 1997).

Figure 1.1

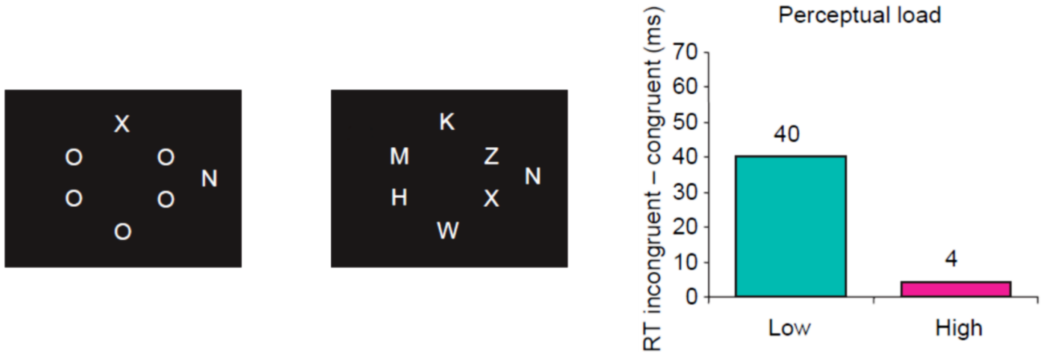

Figure 1.1. Response competition paradigm and typical results for reaction time under low and high perceptual load in the visual search task. In response-competition paradigms participants make speeded responses to a target (here X or N) while ignoring an irrelevant distractor in the periphery (here N). Increases in reaction time for displays with an incongruent distractor (e.g. distractor X and target N) compared to a congruent distractor (e.g. distractor N target N) indicate increased distractor processing. This effect is observed under low load (left panel) but is typically eliminated when perceptual load in the task is high (middle panel).

A natural extension of the observed effects of distractors on reaction time was to directly assess the effect of load on the perception of irrelevant stimuli. The previous measures of reaction time only provided evidence for reduced distractor processing and were not informative about the perceptual consequences of high load, such as changes in sensitivity or conscious perception of irrelevant stimuli. Cartwright-Finch and Lavie (2007) used the *inattentional blindness* paradigm (Mack & Rock, 1998) to measure the effect of load on subjective awareness. They found that reports of awareness for an unexpected irrelevant stimulus presented on the last trial, significantly decreased when participants performed a high load task. This suggests that *inattentional blindness* critically depends on the level of load in the task performed (but see Wolfe, 1999 for an explanation of *inattentional blindness* in terms of “rapid forgetting”). Macdonald and Lavie (2008) adapted the *inattentional blindness* paradigm to directly measure the effect of load on perceptual sensitivity using a dual task. In their experiments participants made a detection judgement, reporting the presence of a critical stimulus that appeared in the periphery on a random half of the trials, while performing a letter-search task at fixation under low or high load. Signal detection analysis of the data allowed for independent measures of perceptual sensitivity and response criterion. These revealed a significant reduction in sensitivity to the peripheral stimulus under high load, while load had no effect on bias. This finding shows a direct effect of load on perception but also demonstrates that load effects can be extended (beyond the previously stated effects on completely *task irrelevant* stimuli) to *task relevant* stimuli in a dual-task setting. When load is modulated in the primary task, while attentional demands in a secondary task remain constant, any reduction in performance in the secondary task may be attributed to resource limitations imposed by the change in load of the primary task.

Reduced sensitivity under high compared to low load was also observed in a study investigating the flicker fusion threshold (Carmel, Saker, Rees & Lavie, 2007). This threshold refers to the speed of a rapidly flickering light at which the light starts to appear continuous and is therefore a measure of an individual’s ability to distinguish rapid changes in light intensity. The results showed a significant reduction in flicker detection sensitivity when participants simultaneously performed a high load letter search task in the periphery (compared to low load). Again, signal detection analysis of the results ruled out a change in response criterion as an explanation for the reduced sensitivity.

Based on the behavioural findings, several neuroimaging experiments have tested predictions derived from Load theory. If high perceptual load in a task indeed reduces distractor processing, this should be reflected in reduced neural activity for irrelevant items. Rees and colleagues (Rees, Frith & Lavie, 1997) were the first to establish that neural responses to irrelevant stimuli are reduced under high perceptual load. The authors presented moving dots in the background while participants made speeded responses to words presented at fixation. In the low load condition participants distinguished a word's case while in the high load condition they monitored its syllables. Motion selective cortical areas (V5, V1 and V2) showed stronger responses evoked by the irrelevant motion under low compared to high load. Similar results have been obtained in a study by Yi and colleagues (Yi, Woodman, Widders, Marois & Chun, 2004) who measured responses to houses and landscapes presented in the background while participants monitored a stream of faces for repetition. The fMRI analysis focused on activation in the parahippocampal place area (PPA), in the medial temporal cortex, which selectively responds to images of places and scenes (Epstein & Kanwisher, 1998). Load was modulated by adding random noise to the face stimuli (a tentative load modulation given the definition of load used above). The results indicated that activation in the PPA was significantly reduced under high load. Moreover, by occasionally repeating the scenes in the background the authors were able to measure changes in response due to adaptation (*fMRI adaptation,* Grill-Spector & Mallach, 2001): Normally, repetition of the same stimuli reduces neural responses when measured with fMRI. However, in this study responses to repeated stimuli were only reduced under low perceptual load in the face task. This result suggests that a distinction between novel and repeated background images was only made when sufficient resources were available. However, when resources were fully engaged in the more demanding face-task, all background images were automatically processed as novel stimuli.

In line with these findings, neural excitability (probed with transcranial magnetic stimulation, TMS) has also been shown to be reduced under high load: While participants performed a letter search at fixation, TMS was administered to the motion sensitive area V5. Stimulation of this area normally produces moving *phosphenes* (brief flashes of light). However, the intensity required to produce such *phosphenes* was significantly greater under high compared to low load in the task at fixation (Muggleton, Lamb, Walsh & Lavie, 2008), indicating a reduction in cortical excitability for areas processing irrelevant distractors. Support for an account of perceptual load in terms of a general reduction in visual excitability came from a study that showed reduced perception under high load even in between stimulus presentations (Carmel, Thorne, Rees & Lavie, 2011). Here participants responded to a critical stimulus that was only occasionally presented in the periphery while they monitored a separate stimulus stream at fixation (looking either for a target defined by colour, in the low load condition, or a target defined by a conjunction of colour and orientation, in the high load condition). Sensitivity to the critical stimulus was significantly reduced under high perceptual load, regardless of whether it was presented simultaneously with the central stimulus or during the empty interval between central stimulus presentations. This finding suggests that the effects of perceptual load are unlikely to be the result of increased biased competition among simultaneously presented stimuli (e.g. Duncan,1980; Torralbo & Beck, 2008) but are instead due to a reduction in visual excitability as a result of reduced baseline cortical excitability in areas processing task irrelevant stimuli (Muggleton et al., 2008).

Together, these results demonstrate effects of load at relatively late stages of the visual processing hierarchy (except for modulations in V5). To what extend these effects depend on modulations at even earlier stages remains to be explored. Several studies have also investigated the question of whether perceptual load can modulate responses at early visual stages. Schwartz and colleagues (Schwartz et al., 2005), presented flashing checkerboard patterns (which evoke strong responses in early visual areas) while participants performed a load task at fixation. The task involved monitoring a stream of small crosses for a single feature (e.g. any red cross) in the low load condition and a specific feature conjunction (e.g. any inverted green cross or upright yellow cross) in the high load condition. The analysis of the imaging data focused on early visual areas (V1-V4) and indeed showed reduced responses under high load in all of these, with the effect of load increasing in the later stages (V2-V4). Furthermore, retinotopic mapping showed that load exerted its strongest effect in the area immediately surrounding the central task, indicating suppressed responses for an annular area surrounding the central task under high load. Such surround suppression has previously been found for spatial attention (Hopf et al., 2006) and is thought to rely on increased recurrent feedback activity in visual cortex (Boehler, Tsotsos, Schoenfeld, Heinze & Hopf, 2009). Furthermore, Hopf and colleagues (2009) only observed surround suppression at the target location when participants had to identify a specific feature conjunction (the orientation of a coloured C) and not when searching for a single ‘pop-out’ feature (a coloured C). It is therefore conceivable that the surround suppression under high perceptual load utilizes the same mechanism.

Others have found similar effects of reduced activity for irrelevant stimuli under high perceptual load in early visual areas, even extending to the lateral geniculate nucleus (LGN; O’Connor, Fukui, Pinsk & Kastner, 2002). Bahrami and colleagues (Bahrami, Lavie & Rees, 2007), extended these findings to unconscious perception by showing that high load reduced responses in early visual cortex to irrelevant stimuli that were presented to one eye and rendered *invisible* by presenting a rapidly changing mask (continuous flash suppression, CFS; Tsuchiya & Koch, 2005) to the other eye. Furthermore, adaptation to unconsciously presented orientations was also reduced under high compared to low load in a foveal task (Bahrami, Carmel, Walsh, Rees & Lavie, 2008).

Overall, the previous findings from behavioural and neuroimaging studies demonstrate reduced neural responses to irrelevant stimuli under high perceptual load throughout the visual hierarchy. However, it remains largely unclear to what extent a simple reduction in neural response strength (or, in other words, signal gain) accounts for the perceptual consequences of high perceptual load. Furthermore, other attentional manipulations, such as spatial or feature cueing, have shown a wide range of effects on neural function at the single cell level, as well as, at the level of the neural population. These include changes in the precision of neural representations in addition to the discussed changes in overall activity and excitability. How perceptual load influences such mechanisms remains to be explored.

1.3 Attentional selection: effects of signal gain vs. tuning

A dominant theme in the attention literature is to distinguish the effects of attention on signal strength from those on selectivity: Does attention improve perception by simply boosting the neural signal representing attended stimuli, or does it change the quality of the signal by modulating the precision of neural responses?

1.3.1 Neurophysiological evidence for attentional effects on gain and tuning

In order to investigate the effect of attention on signal strength, research has focused on visual stimulus contrast as the primary measure. The increase in neural response with increases in stimulus contrast is well described (Albrecht & Hamilton, 1982) and therefore, allows for a direct test of the effects of attention on the earliest components of visual analysis. Furthermore, the relationship between neural firing rate and stimulus contrast (as assessed in monkeys and cats) is mirrored in the behavioural performance of human observers when, for example, detecting a target stimulus of variable contrast. Thus, focussing on rudimentary properties of visual stimuli, such as contrast, allows for a comparison between responses at the single neuron level to that of the neural population determining the behavioural responses of observers. The interaction between neural firing rate and contrast is illustrated by the *contrast-response function* (Figure 1.2). Changes in stimulus contrast are transformed into changes in firing rate by a non-linear (sigmoid) function (Albrecht & Hamilton, 1982; Sclar, Maunsell & Lennie, 1990).

Attention may interact with stimulus contrast to modulate the shape of the contrast-response function in two ways: *contrast gain* and *response gain* (Sclar, Lennie & DePriest, 1989). *Contrast gain*: If attention combines with the response to contrast this could lead to increased contrast sensitivity without changes in the relative firing rate of the neuron. In this case increased sensitivity would be achieved by lowering the threshold at which a neuron responds to a given stimulus contrast. In principle, the attentional modulation would be equivalent to an increase in stimulus contrast (a shift of the contrast-response function to the left). *Response gain*: If attention and contrast contributed independently to a neuron's response, attention would modulate the response to contrast equally across the entire function. In this case, responses would increase proportionally to the original firing rate at a given stimulus contrast without significantly changing the threshold. The characteristic signature of *contrast gain* is retention of the original shape of the contrast response function with attention having the strongest effect at intermediate contrast levels, while the characteristic signature of *response gain* is amplification of the response at high contrast levels.

Figure 1.2

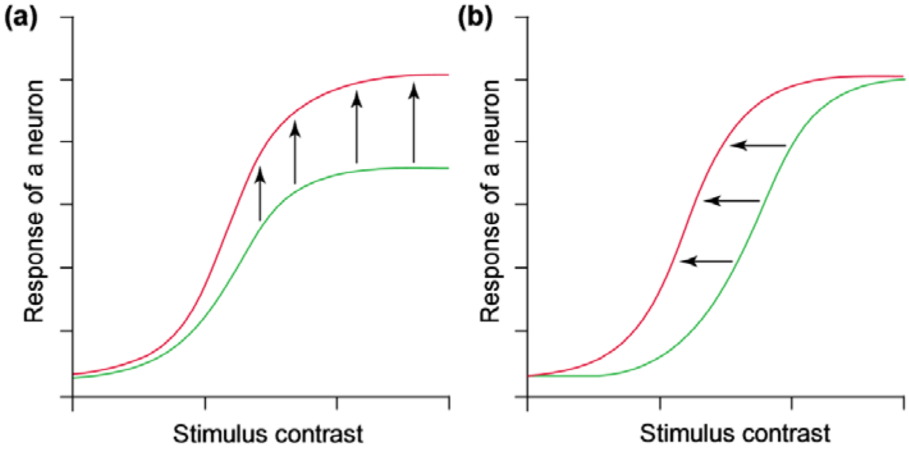

Figure 1.2. Attentional modulation of the contrast-response function. (a) *Response gain*: Attention increases the neural response to contrast by a constant factor across the entire contrast response function. This results in attentional modulation which is strongest at high contrast levels. **(b)** *Contrast gain*: Attention increases neural responses to contrast in a non-linear fashion across the contrast response function, effectively lowering contrast threshold while maintaining the original shape of the function. Attention only modulates responses to contrast at intermediate contrast levels. Adapted from Treue (2001).

Neurophysiological evidence has been obtained in support of both types of gain control. The starting point for such investigation was the finding that sustained attention to a single stimulus presented within a neuron's receptive field increases a neuron's firing rate in much the same manner as an increase in stimulus contrast (Ito & Gilbert, 1999; McAdams & Maunsell, 1999; Motter, 1993; for reviews see Martinez-Trujillo & Treue, 2005; Reynolds, 2005). Reynolds and colleagues (Reynolds, Pasternak & Desimone, 2000) showed that in area V4 of monkey visual cortex, attention caused a left-ward shift in the contrast response function, enabling a neuron to respond to a stimulus at a contrast that previously did not elicit a response. Moreover, the shift was only observed for intermediate contrast levels within the dynamic range of the contrast response function. Attention had no effect at low contrast (10%) and also did not influence firing rate when response was saturated (at 80% contrast). Such a shift at intermediate contrast levels (*contrast gain*) makes a clear prediction for the effect of attention on a neuron's tuning curve.

The tuning curve of a neuron (first formulated by Adrian in 1926, who plotted responses from stretch receptors in muscles as a function of the force applied to the muscle) describes the average firing rate of a neuron in response to a systematic variation of a stimulus parameter. To estimate different parameters of the tuning curve, it is usually fitted with a Gaussian function which quantifies its peak (equivalent to the preferred stimulus feature) and standard deviation (equivalent to the tuning bandwidth). The width of the tuning curve describes a neuron's selectivity to the preferred stimulus property. If attention acts by increasing the effective contrast of a stimulus, this should lead to a multiplicative increase in firing rate across the tuning curve of a neuron, for stimulus contrasts within the neuron's dynamic range. This prediction has been confirmed for orientation tuning curves from neuron's in V4

(McAdams & Maunsell, 1999) and motion direction tuning curves from V5 (Treue & Martinez-Trujillo, 1999). Thus, a contrast gain mechanism (an increase in effective contrast) would explain attentional facilitation when attending to a single stimulus.

Furthermore, the same mechanism can account for the competitive interactions observed when multiple stimuli are presented within a neuron’s receptive field (Reynolds, 2005; Reynolds & Chelazzi, 2004; Reynolds, Chelazzi & Desimone, 1999). Moran and Desimone (1985) were the first to demonstrate that the response of a neuron to two stimuli is predominantly determined by the task relevant stimulus. Neural firing rate increased when monkeys responded to the stimulus with the neuron’s preferred orientation and decreased when they responded to a second stimulus, presented within the same receptive field, with a non-preferred orientation. The finding that firing rates depend on the similarity between a neuron’s sensory preference and the attended stimulus property has been replicated repeatedly (e.g. Chelazzi, Duncan, Miller, & Desimone, 1998; Reynolds & Desimone, 2003; Reynolds et al., 1999; Sheinberg & Logothetis, 2001). Moreover, Reynolds and colleagues (1999) indeed demonstrated how changes in relative contrast of two stimuli presented within a receptive field mirror the effects of attentional modulation. When two stimuli are presented within the same receptive field, one with a preferred orientation and one with a non-preferred orientation, increasing the contrast of the preferred stimulus will increase the overall firing rate, while increasing the contrast of the non-preferred stimulus will reduce the overall firing rate. Thus, the suppressive effect of the non-preferred stimulus is reduced when the contrast of the preferred stimulus is increased. Attention to one of the two stimuli had precisely the same effect as increasing its contrast: the attended stimulus dominated the overall firing rate. The authors came to the conclusion that attention increased the neuron’s contrast sensitivity to the attended stimulus which is equivalent to multiplying its effective contrast (Reynolds et al., 1999). Moreover, such a multiplicative increase in firing rate leaves the neural tuning bandwidth unaffected. This finding together with equivalent results in other visual domains, such as motion (Treue & Martinez-Trujillo, 1999), suggests that attention primarily enhances stimulus salience (amplifying the signal) while leaving other perceived features of the stimulus unchanged (maintaining the same precision).

However, these findings have been challenged by Williford and Maunsell (2006) who showed that attention modulates the contrast response function predominantly through a response gain mechanism (showing attentional modulation even at the asymptotic portion of the contrast response function) or at least a combination of contrast gain and response gain (termed *activity gain*), as well as an additional increase in baseline response. The authors tried to account for this discrepancy between their own and previous findings by pointing out minor differences in their experimental design which may have influenced the monkey's attentional state (in terms of feature based expectations towards low contrast stimuli). Evidence for response gain, however, was also found in a number of other reports (Buracas & Albright, 2009; Huang & Dobkins, 2005; Kim, Grabowecky, Paller, Muthu & Suzuki, 2006), leaving the question of how attention precisely affects the contrast response function open for debate (but see Reynolds & Heeger, 2009 for a possible explanation of the conflicting empirical findings in terms of the relationship between the size of the stimulus and the spatial spread of attention).

Despite the discrepancies between studies investigating the effect of attention on the contrast response function, there has been only a single study that claimed changes in neural selectivity due to narrowing of the neuron's tuning curve (Spitzer, Desimone & Moran, 1988). Since an increase in stimulus contrast has no effect on the width of the tuning curve (this results in a multiplicative change in firing rate across the curve; Sclar & Freeman, 1982; Sclar et al., 1990), any change in tuning width due to attention implies that attention and contrast must rely on qualitatively dissociable neural substrates. Note however, that Spitzer and colleagues (1988) modulated attention in terms of task difficulty and did not alter spatial attention, as did most of the research described above. In their task, monkeys made fine or course discriminations of stimulus orientation or colour. The results showed a narrowing of the neuronal tuning curves for stimuli with the neuron’s preferred orientation when making fine, compared to coarse, discriminations. The results, however, have been revisited by McAdams and Maunsell (1999) who argued convincingly that the findings were misinterpreted due to the definition of tuning width. Spitzer and colleagues (1988) had measured width at a predefined fraction of the height of the tuning curve without taking into account its baseline. However, due to spontaneous activity or response to the non-preferred orientation the baseline is usually non-zero. With a non-zero baseline the width must be measured at a fraction between the base and the peak; thus, preserving width when responses are increased multiplicatively. When the results were re-examined using such an analysis, no change in tuning bandwidth was observed. Therefore, the general consensus that attention does not alter the width of neural tuning curves remained (see also Treue, 2001).

Although a fundamental step towards a complete picture of the neural mechanisms underlying attention, the effects of attention at the single neuron level do not provide a direct link to their perceptual consequences. Human perception is very unlikely to arise from a single neuron (Deneve, Latham & Pouget, 1999; Jazayeri & Movshon, 2006; Pouget, Dayan & Zemel, 2000; 2003). Martinez-Trujillo and Treue (2004), for example, showed that feature based attention may modulate the tuning of the population response in macaque area V5. This study showed that although responses at the single neuron level were modulated by attention in the expected multiplicative fashion (without a change in tuning width), across the population of neurons representing the attended feature the magnitude of individual modulation depended strictly on the similarity between the attended feature (in this case motion direction) and the preferred direction of a given neuron. Responses of neurons that preferred the attended direction were increased, while responses of neurons that preferred the opposite of the attended direction were suppressed (even below their spontaneous firing rate; see Figure 1.3). Here, the combined changes in gain in individual neurons resulted in sharpened tuning for attended features at the neural population level.

Figure 1.3

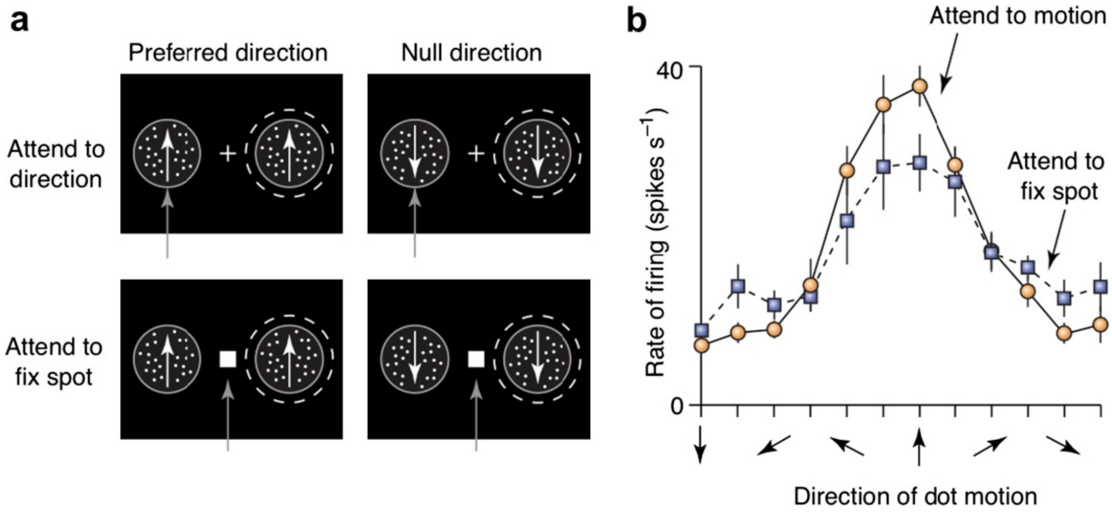

Figure 1.3. The effect of feature-based attention in motion area MT/V5. (a) Neural responses were recorded for two motion stimuli (one presented within the neuron's receptive field (dashed white line) and one outside the receptive field). On some of the trials the

monkeys were cued to attend to the motion presented outside the receptive field (detecting a change in motion speed), on the other trials they detected a luminance change at fixation (grey arrows). **(b)** Firing rate of a representative neuron in MT/V5 was modulated by the attentional modulation. Attending to its preferred direction increased responses, while attending to the null direction decreased responses. Adapted from Martinez-Trujillo & Treue (2004).

This result hints at the complexity of attentional effects on the combined response of the neural population, which has been linked to modulation of human perception in a large number of psychophysical studies, some of which will be highlighted below.

1.3.2 Psychophysical evidence for attentional effects on gain and tuning

In parallel with neurophysiological studies, psychophysical experiments have examined the underlying mechanisms of attention. A large majority of these studies used orientation detection or discrimination tasks. This is due to the fact that just like stimulus contrast, orientation processing has been extensively studied using neurophysiological methods and the findings from neurophysiology and psychophysics have consistently converged (De Valois & De Valois, 1988; Graham, 1989; Ringach, Hawken, & Shapley, 1997). Moreover, performance in orientation discrimination tasks improves with stimulus contrast (e.g. Nachmias, 1967), thus providing a tool to investigate attentional effects on contrast. More recently, orientation perception has been used to probe feature based attention (e.g. Kamatani & Tong, 2006) or estimate attentional effects on population tuning curves (e.g. Baldassi & Verghese, 2005), relying on the findings that fMRI responses increase monotonically with contrast (Boynton, Demb, Glover, Heeger, 1999) and the width of orientation tuning curves is contrast invariant (Sclar & Freeman, 1982).

Behavioural studies have shown that covert spatial attention can improve performance in a variety of ways. Besides explanations in terms of spatial uncertainty or observers' decision criteria (e.g. Palmer, Verghese & Pavel, 2000; Shiu & Pashler, 1994), the current debate has focussed on whether attention primarily acts to increase the signal (*signal enhancement*; e.g. Bashinski & Bacharach, 1980; Cameron, Tai, & Carrasco, 2002; Carrasco, Penpeci-Talgar, & Eckstein, 2000; Carrasco, Williams, & Yeshurun, 2002; Dosher & Lu, 2000a; Ling & Carrasco, 2006a; Lu & Dosher, 1998; Morrone, Denti, & Spinelli, 2002) or if it decreases the response to external noise (*external noise exclusion*; e.g. Dosher & Lu, 2000a, 2000b; Lu & Dosher, 1998; Lu & Dosher, 2000; Lu, Lesmes & Dosher, 2002). Moreover, the two mechanisms are not mutually exclusive and are likely to contribute with different degrees to overall performance, depending on the type of attentional modulation.

The attention literature makes a clear distinction between noise reduction as a mechanism of excluding external noise presented together with the signal (*noise exclusion*) and noise reduction as a mechanism of excluding distractor processing (*distractor suppression*). Both of these mechanisms can best be understood when attention is conceived of as a perceptual filter. In the case of *noise exclusion*, the filter is limited to the location of the attentional focus and its properties may be changed to increase the signal while reducing the noise portion. While for the *distractor suppression* formulation, attention enables the system to specifically select the filter at the attended location, while suppressing information from filters processing distractors outside of the attentional focus. Accounts of spatial attention in terms of *distractor suppression* posit that as spatial uncertainty about the target location or noise from irrelevant distractors increases, performance should be reduced. Spatial uncertainty, for example, increases when the target stimulus is presented alone at very low contrast (close to the observer’s threshold) because the location of the target cannot be distinguished from other (empty) target locations. Increasing the similarity of distractors to the target has a similar effect. The general idea is that increasing spatial uncertainty or the number of distractors reduces performance as the added noise can be confused with the signal. Studies in support of this have argued that pre-cueing the location of a target stimulus should enable the observer to only utilize the filter at the cued location, thus reducing spatial uncertainty (Eckstein, Thomas, Palmer & Shimozaki, 2000; Kinchla, 1992; Nachmias, 2002; Palmer, 1994; Shiu & Pashler, 1994; Sperling & Dosher, 1986). Indeed, others have shown that spatial cueing effects are most pronounced in conditions of low target visibility or low localization accuracy (e.g. Eckstein, Shimozaki, & Abbey, 2002; Morgan, Ward & Castet, 1998), indicating that cueing a location improves performance by simply reducing the number of possible target locations to be monitored (Foley & Schwarz, 1998; Palmer, 1994; Shaw, 1984; Solomon, Lavie & Morgan, 1997) . However, here it is important to distinguish between changes in decision criteria (i.e. the reweighting of information from the target location; e.g. Kinshla, 1992; Shaw, 1984) and suppression of sensory input from distractors.

The effects of spatial uncertainty and decision criteria can be incorporated into experimental design and analysis (e.g. Carrasco et al., 2000; 2002; Eckstein et al., 2002; Morgan et al., 1998) to investigate additional attentional modulation at the level of sensory sensitivity (Eckstein, Peterson, Pham & Droll, 2009). This can be accomplished, for example, by using only a single report location or demonstrating different effects of the magnitude of external noise under different attention conditions. As most uncertainty models explain reduced performance due to increases in set sizes in terms of the increased noise in the display, attention is not expected to affect the impact of such noise. However, dissociations between the effects of attention under low and high noise have been demonstrated repeatedly. For example, the role of *noise exclusion* (noise reduction at the attended location) in attentional facilitation has been extensively documented in a series of experiments by Lu and Dosher where spatial and decision uncertainty was controlled (1998; 2000; Dosher & Lu 2000a, 2000b; Lu et al., 2002). In their perceptual template model, the authors described two primary mechanisms of how attention might improve perception: *stimulus enhancement* and *external noise exclusion*. Both mechanisms are characterized by their distinct effects on performance depending on the level of external noise added to the stimulus. *Stimulus enhancement* improves performance by amplifying the stimulus signal. At high levels of external noise this does not improve performance as both, the signal and the noise, are amplified equally (thus having no effect on discriminability). However, at low levels of external noise, signal amplification is beneficial, as it improves discriminability between the external stimulus and internally generated noise (e.g. from noisy detector responses or loss of information during neural transmission). The alternative to stimulus enhancement, *external noise exclusion*, has the opposite effect. This mechanism relies on better filtering of the signal at the attended location (or feature). It would only be beneficial at high levels of external noise, when there is a sufficient amount of external noise to filter out. Therefore, the model makes clear predictions of how attention would modulate performance (accuracy or d’) at different signal and external noise intensities (Figure 1.4, lower panel). Lu and Dosher concluded from their experiments that in situations where spatial uncertainty for the target is eliminated (for example, by presenting a response cue at the target location), attention primarily improves perception through *external noise exclusion* (Dosher & Lu, 2000a; 2000b; Lu &

Dosher, 2000; Lu et al., 2002). However, *stimulus enhancement* has been observed in some noiseless conditions and in response to peripheral cues (Lu & Dosher, 1998; Lu, Liu & Dosher, 2000).

Figure 1.4

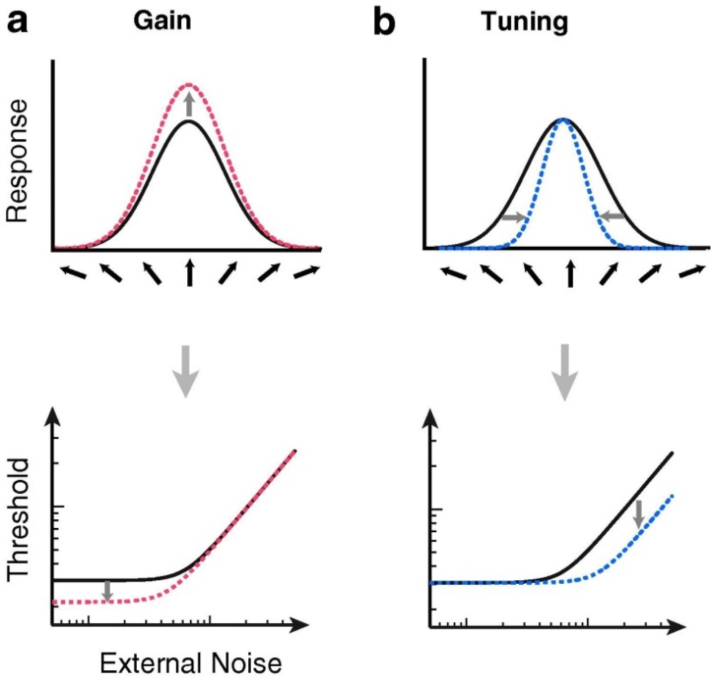

Figure 1.4. Attentional modulations of the perceptual filter (or neural population response) and threshold-versus-noise (TvN) curves. The upper panel illustrates increased gain or sharper tuning of the neural population response, equivalent to signal enhancement or external noise exclusion in the perceptual template model of Lu and Dosher. The lower panel shows the characteristic signature of the two mechanisms on threshold versus noise (TvN) functions (e.g. x-axis: external noise contrast, y-axis: signal contrast at 75% detection accuracy). **(a)** Multiplicative amplification of the neural population response representing the attended stimulus (e.g. a vertical orientation or upwards motion) would result in improved performance only at low levels of external noise. **(b)** Sharpened tuning of the population

response would result in improved performance only at high levels of external noise. Adapted from Ling, Liu and Carrasco (2009).

Others have confirmed that spatial covert attention may operate via *stimulus enhancement* by carefully controlling all variables that could lead to any of the different types of noise reduction mechanisms described above. Carrasco and colleagues (2000), for example, showed that over a wide range of spatial frequencies, contrast thresholds decreased when the target location was preceded by a valid cue compared to a neutral cue. Importantly, the results were obtained under conditions that excluded all variables contributing to external noise reduction (e.g. distractors, masks) while several control experiments ensured that the results were not due to spatial uncertainty. Whether such *stimulus enhancement* actually results in increased perceived contrast at the attended location was subsequently tested and confirmed by Carrasco and colleagues (Carrasco, Ling & Read, 2004).

More recently, these findings have been corroborated by studies that measured the effects of spatial and feature based attention on the perceptual tuning function of the population response (as in Figure 1.4a; Baldassi & Verghese, 2005; Ling, Liu & Carrasco, 2009). Ling and colleagues (2009) varied external noise by modulating global motion coherence for a set of moving dots. They showed that when a central cue was used to direct spatial attention to one of four possible target locations (where observers judged the global direction of moving dots), performance only improved at relatively high levels of coherence (low noise) compared to a neutral cue; the characteristic signature of signal enhancement (increased gain). However, when instead of a spatial cue a feature cue was used to indicate the reference motion direction for a set of moving dots presented at fixation, performance only improved at low levels of coherence (high noise) compared to a neutral cue; the characteristic signature of noise exclusion (sharper tuning). Baldassi and Verghese (2005) also showed an increase in gain across the population response for spatial attention, this time for orientation-selective tuning curves. In contrast, Lee and colleagues (Lee, Itti, Koch & Braun, 1999) reported an increase in gain, as well as, sharpening of orientation tuning curves when an oriented target was attended in isolation, compared to when the target was poorly attended in a dual task setting where observers concurrently performed a highly demanding task at fixation. In both studies (Baldassi & Verghese, 2005; Lee et al., 1999) the authors used a noise masking technique (e.g. Blake & Holopigian, 1985; Legge & Foley, 1980; Ling & Blake, 2009; Ling, Pearson & Blake, 2009; Majaj, Pelli, Kurshan, & Palomares, 2002; Solomon & Pelli, 1994), systematically varying the difference between target and noise orientations. To obtain tuning curves, contrast thresholds were measured for an oriented target at each noise orientation while keeping target orientation constant. However, this paradigm is prone to an “off-channel” looking strategy (Blake & Holopigian, 1985). To illustrate this detection strategy, imagine an observer trying to detect a vertically oriented target. It can be generally assumed that if the target is presented in isolation, observers would utilize those visual filters (or neurons) that respond strongest to the target orientation. However, in the presence of a noise mask containing orientations roughly oriented, say clock-wise from vertical, target detection is enhanced by utilizing filters that are sensitive to a slightly counter-clockwise orientation as this would increase the signal to noise ratio. Thus, any measured reduction in contrast threshold and narrowing of the tuning curve would not correspond to the mechanism in question, the response to the vertical target. Baldassi and Verghese (2005) controlled for this effect by using noise masks with orientations angled symmetrically clock-wise and counter-clockwise from vertical. However, Lee and colleagues did not, which may explain the observed effect on the tuning curves.

The precise conditions under which attention can affect the width of population tuning curves and in that way influence the fidelity of visual perception thus remain to be clarified. Moreover, the effect of perceptual load on tuning curves has not been explored previously and may resolve some of the discrepancies among the studies discussed above.

1.4 Contextual integration and attention

Previous studies on the effects of attention on visual processing have concluded that attentional modulation is strongest in higher processing stages of the visual system and its effects gradually diminish at earlier stages (Maunsell & Cook, 2002). Although in principle current research supports this view, over the past decade it has been repeatedly demonstrated that the magnitude of attentional effects, specifically in early visual areas, greatly depends on the nature of the task and the specific stimulus configuration (e.g. Crist, Li & Gilbert, 2001; Gandhi, Heeger & Boynton, 1999; Ito & Gilbert, 1999; Kastner, De Weerd, Desimone & Ungerleider, 1998; Motter, 1993). Therefore, attentional effects at early visual areas have to be interpreted (and investigated) in light of the specific visual functions subserved by the area in question. In V1, for example, the view has shifted from studying effects on the "classical" receptive field (Hubel & Wiesel, 1962), conceived as a simple feedforward processer of the most basic visual attributes such as location, contrast and orientation, towards the far more intricate effects arising from contextual influences outside the "classical" receptive field. Such "non-classical" properties of neurons have been shown to play a major role in representing complex stimulus configurations even in the earliest cortical visual area, V1 (for review see Gilbert & Li, 2013). This research has demonstrated that early visual areas, in combination with top-down processes and recurrent feedback, can play a functional role in a number of contextual effects ranging from contour integration and surface segmentation to perceptual filling-in and orientation contrast (the tilt-illusion).

1.4.1 Attentional effects on spatial integration

Neurophysiological studies investigating contextual processing have established that a neuron's response to a complex stimulus or scene cannot be predicted just based on the responses to individual parts of the scene presented alone within the receptive field of the cell (Kapadia, Ito, Gilbert & Westheimer, 1995; Kapadia, Westheimer & Gilbert, 1999; 2000). Kapadia and colleagues (1995), for example, showed that responses to a line segment presented to the receptive field of V1 neurons in alert monkeys could be increased considerably (up to 3 times the original firing rate) when a second line segment was presented outside the receptive field. However, presenting a line outside the receptive field in isolation (without a stimulus within the cell's receptive field) had no measurable effect on response. This type of facilitation depends precisely on the closeness of the two lines and their respective orientations. When the lines are collinear, relatively close and have the same orientation, facilitation is most pronounced. Similarly, responses to a line with the neuron's preferred orientation and presented within its receptive field were inhibited when a pattern of randomly oriented lines was presented surrounding the receptive field of the cell. However, when some of the randomly oriented lines were replaced, one by one, with lines presented collinearly and sharing the same orientation of the target, responses from the same neuron increased gradually with an increased number of collinear flankers (Kapadia et al., 1995; Polat, Mizobe, Pettet, Kasamatsu & Norcia, 1998).

Interestingly, in a series of psychophysical experiments, the authors were able to link these findings directly to human perception (Kapadia et al., 1995). Observers were able to reliably detect a line segment at 40% lower contrast when a collinear flanker was present compared to when the line was presented in isolation. Moreover, as previously observed in single cell responses, the effect systematically decreased when the flanking line was moved away from the target along its orientation axis, shifted along the axis orthogonal to its orientation or changed its orientation. This finding has been confirmed by others (e.g. Dresp, 1993; Polat & Sagi, 1993; 1994) and has been observed for luminance judgements of target stimuli presented above their visibility threshold when attention was distributed compared to focussed on the target (Ito & Gilbert, 1999). This latter result suggested that perceptually, a collinear flanker has no effect on a target that is already attended. However, Freeman, Sagi and Driver (2001) challenged this claim by demonstrating strong attentional facilitation from collinear flankers on a central oriented Gabor patch that was equally attended in all conditions. In their study, they used the same stimulus display in two alternative attention conditions where observers either attended to a set of two collinear or two orthogonal flankers, while detecting the presence of the central Gabor patch. Importantly, attention to one set of the flankers had a comparable effect on target threshold as when only this specific set of flankers and the target were presented in isolation. Thus, the effect of each set of flankers in the full stimulus display completely depended on whether they were attended or not (Freeman et al., 2001, 2004; Freeman, Driver, Sagi & Zhaoping, 2003). This and similar results have lent strong support to the view that the perceptual consequences of contextual integration, rooted in the earliest stages of visual analysis, must be interpreted in terms of both, the effects of stimulus properties and comparable effects due to attention.

At the neural level, contextual effects between lines of different orientation have been shown to depend on lateral interactions in early visual cortex (Gilbert & Wiesel, 1989), changes in orientation tuning (Gilbert & Wiesel, 1990) and possibly recurrent feedback from higher cortical areas (due to attention or other top-down factors such as expectation, see Gilbert & Li, 2013). A paradigmatic example of contextual effects in vision is the direct tilt-illusion (also known as orientation contrast). In the classical tilt-illusion, an oriented surround influences orientation perception of a central, oriented stimulus (Figure 1.5).

Figure 1.5

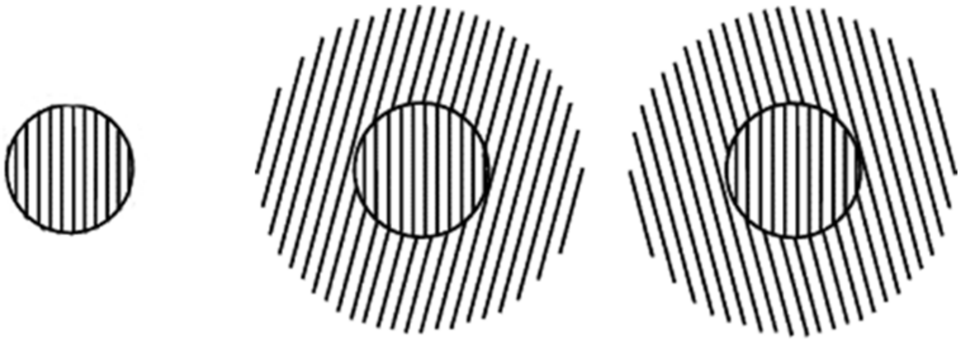

Figure 1.5. Direct tilt-illusion. The same grating appears vertical when presented alone (left) but appears tilted in the opposite direction of the surround (tilt-repulsion) in the context of oriented lines (middle and right).

The tilt-illusion lends itself particularly well to the investigation of attentional effects on contextual influences, as the phenomenon has been extensively studied at the perceptual and neural level. Although the phenomenon itself has been described since the dawn of psychology by Helmholtz, Wundt and others (Luckiesh, 1922), Carpenter and Blakemore (1973) were the first to formally quantify the direct tilt-illusion in detail (i.e. the repulsive shift in perceived orientation that occurs in the presence of a tilted context and does not rely on adaptation). Their analysis suggested that the illusion depends on a mechanism of inhibitory interactions in visual cortex. The authors assumed that the perceived angle depends only on the most active neurons representing the two orientations. As neurons tuned to a certain orientation are known to inhibit responses to similar orientations, the natural peaks of the tuning curves for populations of neurons preferring a specific orientation will be shifted away from the orientation preferred by another population of neurons. If perception is based on the read out from only the most active neurons (the peaks of the population responses) the angle between the two orientations appears expanded. Tolhurst and Thompson (1975) proposed a similar mechanism underlying the illusion and extended the results to stimulus configurations where the inducer surrounds the orientation target (similar to Figure 1.5). Gilbert and Wiesel (1990) confirmed these findings by demonstrating that orientation tuning curves for single cells in cat visual cortex were shifted away from their preferred orientation in the presence of a surround orientation. However, in addition to a shift of the peak of orientation tuning curves the authors predicted that broadening of the orientation bandwidth for neurons preferring orientations away from the surround orientation, as well as sharpened tuning for neurons preferring an orientation close to that of the surround (both were commonly observed in their single cell recordings) should both result in a repulsive shift of roughly the same magnitude as that observed in psychophysical studies. Similarly, Felsen and colleagues (Felsen, Touryan, & Dan, 2005), showed that tuning curves of neurons in cat primary visual cortex were shifted away from their preferred orientation in the presence of spatial (i.e. tilt-illusion) or temporal context (i.e. tilt after-effect) and showed that this could lead to a reduction in redundant signals in the population response, increasing the efficiency of visual representation. Others have confirmed that local inhibition and mechanistic changes to tuning curves are likely to underlie the perceived tilt-illusion (e.g. Felsen, Shen, Yao, Spor, Li & Dan, 2002; Li, Thier, & Wehrhahn, 2000; Müller, Metha, Krauskopf & Lennie, 2003; Sengpiel, Sen, & Blakemore, 1997).

However, how attention affects these processes has rarely been explored. Sundberg and colleagues (Sundberg, Mitchell & Reynolds, 2009) investigated attentional effects of contextual modulation by measuring single cell responses to two stimuli presented together, one within the cell's "classical" receptive field and one outside the receptive field. The presence of the stimulus in the surround suppressed responses to the stimulus within the receptive field. However, suppression increased when attention was directed to the surround and decreased with attention focused on the centre stimulus (see also Moran & Desimone, 1985). This finding supports the notion that spatial attention not only increases the signal of the attended stimulus but also reduces the influence of task-irrelevant items in the surround. Applied to the phenomenon of the direct tilt-illusion, the results predict attentional modulation of centre surround interactions depending on attentional allocation to a stimulus or its context.

Although a number of studies have investigated perceptual effects of attention on the tilt after effect (TAE), the temporal equivalent to the direct tilt illusion, where adaptation to an orientation results in the perceived repulsive tilt of a subsequent test stimulus (e.g. Liu, Larsson, & Carrasco, 2007, Spivey & Spirn, 2000), there is very little behavioural data on attentional effects on the direct tilt illusion (Schwartz & Coen-Cagli, 2013). However, this has been examined for a related phenomenon in visual motion. When a central patch of coherently moving dots is surrounded by a ring of motion in a different direction, the central motion is misperceived as repulsed away from its original direction (the analogue to the tilt-illusion in the motion domain). Tzvetanov and colleagues (Tzvetanov, Womelsdorf, Niebergall & Treue, 2006) demonstrated that the strength of this illusion increased with attention to the surround and suggested that attention to the central target alone may reduce the bias. These findings are in line with a computational model that predicts an increase in the strength of the illusion when centre and surround were both attended compared to when attention was directed towards the centre alone (Schwartz & Coen-Cagli, 2013). Mareschal and colleagues (Mareschal, Solomon & Morgan, 2009) reported that when an oriented target was presented together with two oriented flankers in the periphery, attention directed to the flankers reduced the perceived tilt-illusion of the target. However, this did not improve performance in identifying the target orientation (limited by crowding). Thus, spatial attention to the flankers reduced the tilt-illusion without improving visual acuity, suggesting an earlier locus for the effects of crowding compared to contextual modulations.

How such illusions are influenced by withdrawing attention from the entire stimulus (including target and inducer) has not been explored either. However, the findings discussed above raise the possibility that perceptual load may influence contextual interactions among unattended stimuli and their context, for example, if perceptual load modulates the strength of lateral inhibition among orientation detectors.

1.5 Temporal integration and attention

Another formulation of contextual influences modulated by attention can be found in *visual masking*. Masking in its most general form, occurs whenever perception of a stimulus is reduced or eliminated by the presence of a second stimulus (Breitmeyer & Ogmen, 2000). Therefore, it extends to the contextual modulations discussed above, and has provided an essential tool in psychophysics to investigate a large variety of spatial visual properties. Moreover, by presenting a mask *after* a target stimulus, the temporal dynamics of spatial integration can be investigated. Impaired perception of a target stimulus due to the presentation of a mask that temporally follows the target stimulus is called *backward masking*. This phenomenon, interesting in its own right due to the counterintuitive observation that a mask can reduce perception of a target that is presented first (Bachmann, 1994; Breitmeyer, 1984), has been used extensively to probe the temporal order of visual information processing in the psychophysical (e.g. Bachmann, 1984; Bowen & Wilson, 1994; Michaels & Turvey, 1979; Muise, LeBlanc, Lavoie & Arsenault, 1991) and neurophysiological domains (e.g. Keysers & Perett, 2002; Kovacs, Vogels & Orban, 1995; Macknik & Livingstone, 1998; Rolls Tovee & Panzeri, 1999; Thompson & Schall, 1999).

Furthermore, it has been proposed that some of the neurophysiological mechanisms required for contextual effects (discussed above) are impaired by curtailing the time needed for neural integration to produce a stable percept (e.g. cortical re-entrant processing; Lamme, Zipser & Spekreijse, 2002). Thus, if the limited perceptual resources under conditions of high perceptual load affect the processes involved in spatial contextual integration, these are likely to modulate the temporal profile of target and mask integration. Moreover, the temporal locus of such modulation may provide valuable clues as to the levels of visual processing affected by load.

1.5.1 Attentional modulation of backward masking

There is considerable evidence indicating a role for attention in modulating masking curves obtained with backward masks. Some have used indirect measures to assess the effect of attentional mechanisms. Ramachandran and Cobb (1995), for example, showed that perceptually grouping a target stimulus (that was subsequently masked) with another stimulus (that was not masked) reduced the effectiveness of backward masking, compared to when participants were instructed to group two other nearby stimuli (that did not include the masked target). Others have shown that when targets consist of stimuli that are highly relevant to the observer (such as a person's name) metacontrast masking (discussed below) appears to be reduced compared to when physically similar control targets are used (Shelley-Tremblay & Mack, 1999); suggesting a role for top-down influences on perception of masked stimuli. This finding is supported by a large number of studies showing masked priming effects where two stimuli are presented in succession and the response to the second stimulus is enhanced if the first stimulus is semantically related to the second, even when it is rendered invisible by a mask (e.g. Cheesman & Merikle, 1986; Debner & Jacoby, 1994). Naccache and colleagues (Naccache, Blandin & Dehaene, 2002) showed that this effect depends on the temporal allocation of attention to the period where the subliminal prime is presented and does not occur without attention.

More direct investigations of attentional effects on backward masking were conducted by altering the number of target elements in the display. Early studies suggested that focused attention on the target reduced the effectiveness of a mask compared to spatially distributed attention (e.g. Averbach & Coriell, 1961). In their study, participants either identified a single element at a known location (followed immediately by a backward mask at different SOAs) or they identified an element embedded in a multi-element target display, whose location was only known once the mask appeared. Thus, attention could be allocated to the target location before the target-mask sequence in the single element condition but not in the multi-element condition. The authors concluded that the more diffuse attentional state and delay in attentional allocation to the target location in the multi-element display led to delayed target processing as well as reduced perceptual sensitivity of the target, resulting in stronger masking. Others provided supporting evidence for this conclusion (Spencer, 1969; Spencer & Shuntich, 1970) by comparing target identification accuracy at various SOAs in a single-letter and 12-letter target array, similar to Averbach and Coriell (1961). However, in addition performance was tested in both conditions at three different levels of mask intensity. With both displays, mask intensity increased the effectiveness of the mask. However, in the single-letter display masking was strongest when target and mask were presented simultaneously (SOA = 0ms) and recovered for all mask intensities by ~150ms, while in the 12-letter display masking strength further increased at SOAs greater than 0ms and did not fully recover until ~300ms SOA.

Michaels and Turvey (1979) proposed a two-stage model of backward masking inspired by the results from set size manipulations, where the early effects of the mask (at short SOAs) were attributed to sensory integration between target and mask or sensory suppression of the target by the mask, while masking effects observed at longer SOAs (e.g. Enns & Di Lollo, 1997; Spencer & Shuntich, 1970), where attributed to higher level, attentional interactions. More specifically, the mask was thought to interrupt attentional processes devoted to transferring information from iconic visual memory to non-visual forms of memory (e.g. working memory or verbal short term memory) which were thought to be unaffected by visual masks (e.g. Neisser, 1967; Sperling 1960; 1963; 1967). Michaels and Turvey (1979) supported their theory of *interruption masking* by showing that semantically meaningful information was less prone to suppression by backward masking compared to meaningless information, supporting the claim that masking acted on the readout from iconic memory which was assumed to be more efficient for higher order semantic information. Others have reported similar findings of reduced masking effectiveness for target stimuli that are known to be more resistant to perceptual degradation in *inattentional blindness*, such as familiar names or faces (Shelley-Tremblay & Mack, 1999).

More recently, others have used very similar paradigms, for the most part replicating the previous findings on set size and masking strength (e.g. Enns & Di Lollo, 1997; Tata, 2002; Tata & Giashi, 2004). Enns & Di Lollo (1997) for example, presented a target and mask either at a single, known location or at one of three possible target locations (their Experiment 1). Tata (2002) used a spatial cue to indicate target location in single or multi-element displays. Both studies found that masking was more effective when the location of target and mask was unknown or the set size of elements in the display (possible targets or target locations) was increased. These results were interpreted in terms of *object substitution* (Di Lollo, Enns & Rensink, 2000; Enns, 2004; Enns & Di Lollo, 1997; discussed in more detail in the next section), a theory similar in spirit to *interruption theories* of masking (e.g. Michaels & Turvey, 1979).

However, a critical shortcoming of masking experiments involving changes in set size is their limited explanatory scope with regard to the underlying attentional mechanisms. As discussed previously the behavioural benefits with small set sizes can be well explained in terms of spatial and decision uncertainty (Baldassi & Burr, 2000; Cameron et al., 2004; Kinchla, 1992; Morgan et al., 1998; Shiu & Pashler, 1994; Solomon et al., 1997). This is particularly the case when the target is embedded in the search display as one of the elements (and is not spatially separated or presented at a known location). Accordingly, increasing the number of distractors (or empty target locations), increases the amount of noise that can interfere with identification of the target signal.

In case of any attentional involvement in set size manipulations, it is likely to modulate the interaction of target and mask at a number of different stages in the visual system. Recent neurophysiological evidence can clarify how the *interruption* models of backward masking discussed above would be implemented in the brain. Kovacs and colleagues (1995), for example measured neural responses to masked stimuli in monkey pretemporal cortex (IT), which is likely to underlie shape recognition processes (Gross, Rocha-Miranda & Bender, 1972). The authors showed that presentation of a backward mask reduced the number of spikes in neurons responsive to the target, effectively reducing a neuron’s capacity to selectively respond to a distinct shape. Thus, under those conditions where temporal integration is required for the system to produce reliable behavioural responses, the presentation of the mask can interrupt such integration even at higher levels in the visual system. This finding has been corroborated and extended to activity in early visual areas by experiments investigating masking effects on figure ground segregation. Lamme and colleagues (Lamme, Rodriguez-Rodriguez & Spekreijse, 1999) first established two different components of responses from cells in V1 when presented with simple figure ground stimuli comprising line segments tilted in one of two orientations. An early component showed no difference in neural response when cells sampled the figure or the ground. However, a late component of enhanced response (starting around 100ms after stimulus onset) was only observed when cells sampled the figure but not when they responded to the background. In a follow up study, Lamme et al. (2002) then demonstrated that the early responses in V1, (corresponding to simple detection of line orientations) were unaffected by a pattern backward mask (the mask had no effect on recorded neural responses at short SOAs), while the later responses, involved in contextual integration and figure ground segregation, were reduced by the presentation of the mask (at SOAs beyond 100ms). As only the latter of these processes is thought to rely on recurrent feedback, the authors interpreted their results in favour of *interruption* or *object substitution* theories of backward masking, arguing that the mask interrupted late feedback from higher visual areas to V1 (Lamme, 2000; Lamme & Roelfsema, 2000). According to this account, backward masking occurs specifically because of a mismatch of information between higher and lower levels in visual processing: While higher areas feedback information about the target stimulus to lower visual areas, here processing of the target has already been replaced by that of the mask, interrupting further processing of target stimulus information in higher levels.

1.5.2 Types of backward masking and relation to theoretical models

Backward masking discussed in its general form above, can be subdivided into a number of different categories (e.g. Breitmeyer & Ogmen, 2000) depending on shared spatial and featural properties between target and mask. For the current thesis the focus will be on three types of backward masking defined by the properties of the mask in relation to those of the target:

1) *Pattern masking*: here the mask is similar to the target in terms of its visual structure, such as contours, and is presented at the same location as the target.

2) *Metacontrast masking*: in this case the mask is presented closely surrounding, but not overlapping, the target location. Thus, the mask shares some of its contour features with those of the target.

3) *Object substitution masking*: this type of masking uses a mask which is both, dissimilar in structure and even further removed from the target location.

While attentional effects on *pattern masking* and *metacontrast masking* have, for the most part, only been observed with indirect measures (Ramachandran & Cobb, 1995; Shelley-Tremblay & Mack, 1999; but see Tata, 2002, discussed below) *object substitution masking* appears to be strongly determined by focal attention (Di Lollo et al., 2000; Enns & Di Lollo, 1997). The effects of backward masking on target perception are measured by acquiring performance indicators, such as target detection accuracy or d', at various stimulus-onset-asynchronies (SOAs) between target and mask. The specific shape of the obtained masking curve, plotting performance as a function of SOA, is predicted to differ for different types of masking and types of models aimed at explaining the underlying mechanism of target-mask integration.

For *pattern masking* (1), two primary mechanisms have been proposed that independently interfere with perception of the target. The first mechanism is simply based on the inability of the visual system to distinguish the target from the mask due to its poor temporal resolution at short SOAs (up to ~100ms). In this case the masking effect is similar to presenting mask and target simultaneously (effectively increasing the level of noise in the stimulus) and is thought to depend on the integration of early visual representations, commonly referred to as *integration masking* (Breitmeyer, 1984; Kahneman, 1968; Scheerer, 1973; Turvey, 1973). The characteristic masking curve as a result of *integration masking* is U-shaped and centred around the 0ms SOA (when target and mask are presented together; *integration masking* can also be effective at negative SOAs, when the mask is presented before the target). A second mechanism, *interruption masking* (already discussed above) is thought to be prevalent at later SOAs, where the mask interferes with further processing of the target at the stage of object recognition (Bachmann & Allik, 1976; Michaels & Turvey, 1979; Scheerer, 1973). Thus, a typical masking curve that shows *interruption masking* effects has its lowest point (where masking is strongest) beyond the 0ms SOA mark. This division is supported by the finding that changes of the physical characteristics of the mask, such as an increase in contrast, alter *integration masking* and affect the masking curve at short SOAs but do not affect *interruption masking* at later SOAs (Breitmeyer, 1984; Scheerer, 1973; Spencer & Shuntich, 1970). In contrast, changing the layout of the display, for example, by increasing the number of possible target locations, significantly increases *interruption masking* while having no effect on *integration masking* (Spencer & Shuntich, 1970).

In *metacontrast masking* (2), the mask tightly surrounds the target location without overlapping it. Unlike pattern masking, this type of mask does not reduce target visibility when presented together with the target. Therefore, metacontrast masking cannot be explained by degradation of the target signal due to added noise from the mask (i.e., *integration masking*). Also, when the mask is presented first, followed by the target, or if the mask is presented at late SOAs (typically beyond 150ms with focussed attention on the target-mask sequence), no masking occurs. However, as metacontrast masking is most pronounced at very short SOAs around 50-

100ms, it has generally been assumed to rely on local mechanisms in early visual cortex related to contour integration. Challenging classical *integration* accounts of masking but not relying on immediate sensory integration to explain the effects of metacontrast, Breitmeyer and Ganz (1976; Breitmeyer, 1984) proposed that metacontrast masking relies on interference between two channels carrying distinct types of visual information: one channel rapidly conveys information about transient changes such as target onsets and offsets while the other channel conveys information more slowly, about other stimulus features such as colour or shape. Metacontrast masking is thought to occur when the transient activity elicited by the mask onset inhibits the sustained response from the target in the slower channel. This and similar *dual-channel* models suggesting local contour interactions (for review see, Breitmeyer, 1984) have received support from studies showing that metacontrast masking strictly depends on the proximity between target and mask, with masking strength greatly decreasing when target and mask are separated by more than half a degree of visual angle (Breitmeyer, 1984; Growney, Weisstein & Cox, 1977).

However, this view of metacontrast masking is challenged by a parallel line of research investigating the effects of higher level cognitive processes (e.g. Averbach and Coriell, 1961; Enns & Di Lollo, 1997; Michaels & Turvey, 1979) that showed how under some conditions metacontrast masking was not effective when a target was presented at a known location before presentation of a metacontrast mask but when the target was part of an array of letters with the mask presented only at the location of the target, the typical U-shaped, metacontrast masking function was observed. As already discussed, others have reported similar results that cannot explain the effect of metacontrast masking in terms of local contour inhibition alone (Michaels & Turvey, 1979; Uttal, 1970), showing, for example, that masking strength was modulated by similarity of the geometric shapes of target and mask.

More recently, the interpretation of the findings of set size manipulations in terms of spatial attention (Enns & Di Lollo, 1997; Ramachandran & Cobb, 1995; Shelley-Tremblay & Mack, 1999; Tata, 2002) has received further support from experiments using masking techniques that reduce the influence of local target-mask interactions even further. The main idea is that masking can be reduced or even prevented when focused attention on the target strengthens the target's representation. Di Lollo and colleagues (Di Lollo et al., 2000) tackled the question of attentional influences in backward masking by employing a *common-onset* paradigm (Di Lollo, Bischof, & Dixon, 1993) where target and mask are presented simultaneously and only the duration of the mask is varied. This paradigm rules out any effects in terms of inhibitory mechanisms depending on the onsets of target and mask contours (as in the *dual-channel* model discussed above; Breitmeyer & Ganz, 1976). The authors showed that masking strength increased with longer mask durations, and this effect was significantly modulated by spatial cueing of the target, target pop-out or increased set size (the number of possible target locations). Moreover, the results were generalized to masking with four small dots (instead of the contour mask surrounding the target location used in classical metacontrast masking) to account for any other low level inhibitory effects. The results with this type of mask (*object substitution masking* (3)) showed a similar effect of set size. The sensitivity of this type of masking to spatiotemporal deployment of attention was interpreted in terms of neural re-entrant processing, where stimulus information is fed back to earlier levels of analysis to be compared with new input in order to maintain a stable percept which then gains access to conscious awareness (Di Lollo et al., 2000; Lamme et al., 2002). However, Di Lollo et al.’s (2000) findings and theoretical interpretation critically depend on the processes associated with the prolonged presentation of the mask in the *common-onset* paradigm. In classical metacontrast masking, however, the mask is always presented for exactly the same length as the target. To bridge this gap and specifically investigate attentional effects in metacontrast masking, Tata (2002) measured target detection with spatial cueing, target pop-out and set size variations, while keeping target and mask durations constant. He found that the attentional manipulation in each of these experiments had a strong effect on masking strength.

However, the majority of the studies discussed above are limited in their extent to explain the underlying attentional factors. For example, it is important to note that in all experiments in Tata’s study (2002), as well as in Di Lollo et al. (2000), attention could only be focused on the target location once the mask occurred (since the mask always acted as a cue indicating which of the stimuli in the display to respond to). Therefore, it is very likely that the effectiveness of masking in these studies depends not only on attentional selection of the target but also, and presumably to a greater extent, on selective attention to the mask. In classical metacontrast masking (e.g. Tata, 2002) the sudden onset of the mask should effectively capture attention (e.g. Jonides & Yantis, 1988; Yantis, 1993) whereas, in *object substitution masking* using the *common onset* paradigm (e.g. Di Lollo et al., 2000) the prolonged presence of the mask in isolation significantly increases its salience. Whether backward masking is affected by attention in situations where target location is fixed remains to be explored. In one recent study, effects of spatial attentional on metacontrast masking were investigated with endogenous cues at fixation (Boyer & Ro, 2007), directing attention to a valid or invalid target location. Valid cueing reduced masking strength compared to invalid cues and recovery from masking was faster with valid cues. This finding suggests that spatial attention may indeed influence metacontrast masking when effects of spatial uncertainty are accounted for (discussed in the General Discussion). However, it remains unclear how limiting available perceptual resources will influence target and mask interaction and whether this would affect target processing at early or later stages in the visual system.

1.6 General methodological approach and overview (Chapters 2 to 5)

The aim of the current thesis is to investigate the mediating mechanisms that underlie the perceptual effects of limiting attentional resources through perceptual load. To probe mechanisms of perceptual load at early levels of visual analysis and quantify their relationship to perceptual experiences, several psychophysical methodologies were employed.

Throughout the empirical chapters, perceptual load was always modulated in a separate, concurrent visual search task surrounding a target stimulus at fixation. Processing of the target stimulus under low and high perceptual load was quantified by recording target detection and discrimination accuracy scores, using signal detection theory to estimate sensitivity and response criteria, and estimating detection thresholds and points of subjective equality using psychophysical staircase procedures.

In Experiments 1 to 3, psychophysical orientation tuning curves were obtained to establish the mechanism by which high perceptual load in a task surrounding fixation reduces orientation perception of stimuli in a concurrent task at fixation. The results showed a reduction in gain when attentional resources were withdrawn from the target stimulus under high load, which has typically been observed in studies modulating spatial attention. However, in addition to a change in gain, high perceptual load modulated the selectivity of the population response by broadening the width of the psychophysical orientation tuning curve. This result critically depended on the level of perceptual load in the task and not on the physical properties of the stimuli or a reduction of effective contrast at the target location (Experiments 2 and 3).

Having established that perceptual load alters orientation selectivity, experiments 4-7 investigated how this affects contextual integration in the orientation domain, specifically in the direct tilt-illusion. The results suggest that perceptual load in an unrelated task modulates the strength of the perceived tilt-illusion. However, the sign of the modulation depended on the spatial configuration of the stimuli, as well as, on whether the context inducing the tilt-illusion was consciously perceived. Experiments 6 and 7 confirmed that the previously observed increase in the perceived tilt-illusion under high load was not due to the specific stimulus parameters of orientation target and context, but instead depended on the spatial separation between target and context and the unconscious presentation of the context. These results support a role for perceptual load in modulating lateral interactions at the neural population level or altering active suppression of task irrelevant stimulus features.

Neural processes involved in contextual effects in vision are known to occur at later points in time and extend over longer periods, even in early visual areas, compared to processes representing single stimulus features. Experiments 8 through 13, therefore, extended the findings from the spatial to the temporal domain by investigating the effects of perceptual load on different types of backward masking, exploiting temporal signatures of modulations by load to estimate their locus in the visual hierarchy.

Chapter 2

The effect of perceptual load on orientation processing: Gain versus Tuning

2.1 Chapter Introduction

As detailed in the general introduction, the effects of attention on visual perception depend on the level of perceptual load in the task (Lavie, 2005). Due to the capacity limits of visual perception, tasks involving higher perceptual load (e.g. search tasks involving many similar items, or tasks requiring complex perceptual discriminations; e.g. Lavie, 1995; Lavie & Cox, 1997) result in reduced visual cortex responses to unattended stimuli (e.g. Rees et al., 1997; Schwartz et al., 2005; Yi et al., 2004; see Lavie, 2005, 2010 for reviews) and lead to the experience of *inattentional blindness* (Carmel et al., 2011; Cartwright-Finch & Lavie, 2007; Macdonald & Lavie, 2008; Simons & Chabris, 1999). The effects of perceptual load on neural responses have been established across different load manipulations and in a variety of tasks. These effects are found to extend throughout the visual cortical hierarchy, from occipital cortex including primary visual cortex area V1, the superior colliculus and LGN (e.g. Bahrami et al., 2007; O'Connor et al., 2002; Rees et al., 1997; Schwartz et al., 2005), through to cortical areas involved in the perception and recognition of complex images, meaningful objects and scenes (e.g. Pinsk, Doniger, & Kastner, 2004; Yi, et al., 2004).

However, a simple reduction in neural response signal, although certainly affecting our visual experience, may only be one part of the explanation. According to signal detection theory, for example, successful visual detection and discrimination do not only depend on the strength of the signal (in other words the signal gain) but also on the extent to which the signal is precisely tuned (Green & Swets, 1966). Thus, reduced perception, as well as, reduced neural response under high load may be the result of either reduction in the signal gain or broadening of the tuning profile of the population representations for a given stimulus, or both.

In the experiments in this chapter, I therefore examined whether reduced perception under load is due to a reduction in gain or broadened tuning or a combination of the two. The experiments focus on perception of orientation, an elementary operation of visual perception, and assessed the effects of perceptual load on psychophysical measures of orientation tuning curves obtained by using a noise-masking paradigm (Baldassi & Verghese, 2005; Blake & Holopigian, 1985; Legge & Foley, 1980; Ling & Blake, 2009). Orientation perception is ideally suited to assess the effects of perceptual load on early visual responses as the psychophysical tuning curves provide a link to the responses of the underlying population. Indeed, in many cases the characteristics of tuning curves recorded from single cells are strikingly similar to those obtained from human observers in behavioural tasks (for review see Neri & Levi, 2006), indicating that such psychophysical measures are particularly useful in revealing the underlying neurophysiological substrate.

2.2 Experiment 1

The purpose of Experiment 1 was to establish the effect of perceptual load on orientation tuning. Observers performed a visual search task with either low or high perceptual load, while also detecting an oriented stimulus embedded within a noise mask (Figure 2.1a). The orientation content of the mask was varied, while the orientation of the target remained constant. Contrast detection thresholds were then obtained for each noise orientation in order to plot orientation tuning curves. The width of the curves indicates the precision of the target signal (tuning) while amplitude and overall vertical shift indicate the strength of the signal (gain).

2.2.1 Method

Observers. Twelve observers, six of whom were female (aged 20-28, *Median* = 25), participated in the first experiment. All were recruited from the University College London subject pool, had normal or corrected-to-normal visual acuity and were naïve as to the purpose of the study. All aspects of the study were in accordance with the local ethics committee at University College London.

Apparatus & Stimuli. Stimuli were created using MatLab (2007a, The MathWorks, Nattick, MA) and Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) and presented on a 21” Monitor (1024 x 768 pixel resolution, 75Hz refresh rate) in a darkened room. Viewing distance was maintained at 57cm with a chinrest.

In all experiments reported in this chapter, observers indicated the location of a vertically oriented Gabor patch (1.5° of visual angle in diameter, with a spatial frequency of 6 cycles-per-degree, cpd) which appeared with equal likelihood in one of two locations centred at 1° above or 1° below fixation. The target Gabor patch was embedded within a round noise-mask (15% root-mean square, rms, contrast, 5° in diameter) presented at the centre of the screen. The noise was low-pass spatial frequency filtered (10 cpd cut-off) and band-pass filtered in the orientation domain with a 20° bandwidth. The centre frequency of the noise had one of seven possible orientations in each block: 0°, 8°, 16°, 24°, 32°, 40° or 90° from vertical. Figure 2.1c illustrates how detection of the vertical Gabor patch becomes increasingly difficult as noise orientations become more similar to the target orientation (i.e. from right, 90°, to left, 0° mean noise orientation). Noise band-pass orientations were angled symmetrically clockwise and counter-clockwise from vertical in order to prevent off-channel looking (e.g. Blake & Holopigian, 1985; see Chapter 1).

A letter search task was used to vary the level of perceptual load (Figure 2.1a). Eight, dark grey letters (subtending 0.6° x 0.9° of visual angle) were presented together with the central noise-mask and target Gabor, equidistant from fixation centred at 4.5° eccentricity. Observers searched for a target letter (Z or N) among either heterogeneous non-target letters (randomly selected from X, E, K, L, H, M, F or T; high load condition) or homogenous non-target letters (all Vs; low load condition). Varying the similarity between distractors and target in this way is a well-established perceptual load manipulation that has been shown to effectively reduce or eliminate distractor interference (e.g. Lavie & Cox, 1997) and reduce sensitivity (e.g. as measured with d'; Macdonald & Lavie, 2008) to additional stimuli presented in the periphery in dual-task settings (e.g. Carmel et al., 2011; Macdonald & Lavie, 2008). Importantly, the effects of perceptual load in visual search have been dissociated from general effects of task difficulty, pointing to search efficiency as one of the determining factors of perceptual load (Lavie & Cox, 1997; Roper, Cosman, & Vecera, 2013).

Figure 2.1

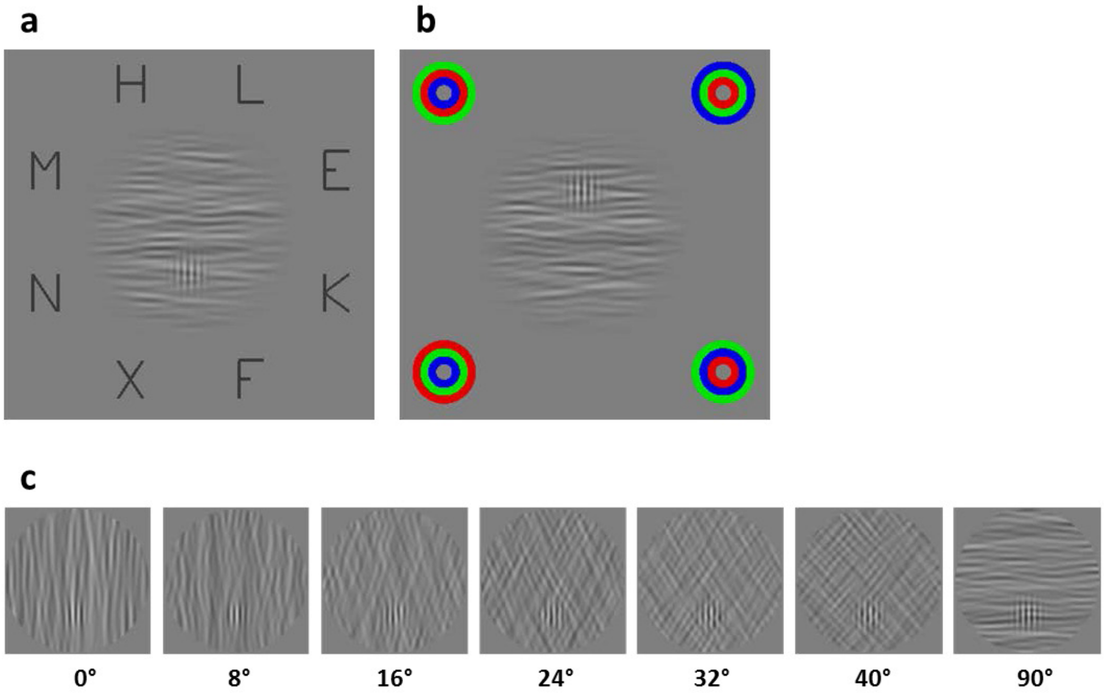

Figure 2.1. Schematic illustration of stimulus displays and noise masking procedure in the high load conditions of Experiment 1 and 2. (a) In Experiment 1 observers performed a dual-task, searching for a target letter (Z or N) among either heterogeneous non-target letters (high load condition, shown) or homogenous non-target letters (all Vs, low load condition, not shown) while detecting the location of a small Gabor probe presented in orientation band-passed noise. **(b)** In Experiment 2 observers performed a colour-conjunction search task instead of the letter search task. The target was always one of two pre-defined, unique colour conjunctions and appeared either among heterogeneous non-target discs (high load condition, shown) or among homogeneous discs (low load condition, not shown). **(c)** The central frequency of the orientation band-passed noise varied from being identical in orientation (0°) to being orthogonal in orientation (90°) to that of the target. The figure (**c**) illustrates how detecting the vertical target becomes increasingly difficult when viewing from right to left. The target orientation remained constant throughout the experiment at 0° (vertical).

Procedure. At the beginning of each trial a fixation cross was presented at the centre of the screen for 1200ms, followed by simultaneous presentation of all stimuli for 160ms (Figure 2.2). Subsequently, participants responded first to the letter search (indicating the target letter by pressing one of two designated keys on the left side of the keyboard) and then to the orientation discrimination task, indicating whether the Gabor patch had been detected above or below fixation (using the up and down arrow keys on the right side of the keyboard).

Figure 2.2

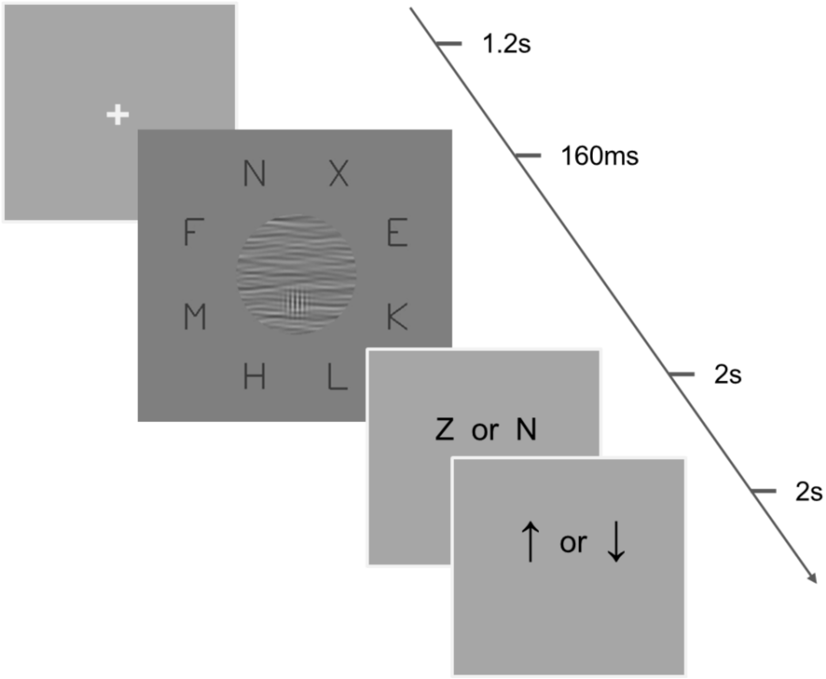

Figure 2.2. Schematic illustration of trial sequence and stimulus duration in Experiments 1 and 2. The load task (letter search or colour conjunction search, not shown) was presented together with the orientation detection task at fixation. Observers always responded first to the visual search target and then indicated whether the oriented target stimulus had been presented in the upper or lower half of the noise patch.

The short stimulus duration of 160ms was chosen to preclude any eye movements. Furthermore, observers were instructed to prioritize the letter search task over the orientation detection task and respond to both as quickly and accurately as possible. Note that the orientation of the Gabor patch was fixed at 0° (vertical) across all trials. Noise band-pass orientations (0° to 90°) and load levels (high vs. low) were blocked and their order counterbalanced across participants and sessions. Observers completed 4 sessions of 14 blocks each (resulting in 4 contrast thresholds obtained for each load condition and noise orientation, with 40 trials per threshold estimate). Different sessions were conducted on separate days and observers rested for two minutes after each block.

Psychophysics. In all experiments of Chapter 2, an adaptive staircase procedure (QUEST; Watson & Pelli, 1983) was employed to estimate contrast thresholds at 75% accuracy for the target Gabor patch at each of the seven noise band-pass orientations. Staircases for each noise orientation and load condition were blocked and their order pseudo-randomized and counterbalanced across sessions and observers. In each of four sessions, a single threshold was obtained for each of the seven noise band-pass orientations under high and low perceptual load.

For every observer, Gaussian functions were fit to the data obtained from averaging the four contrast thresholds measured at each noise level. The functions were centred on 0° and assumed to be mirror-symmetric. Subsequently, bandwidth (full width at half maximum, FWHM) parameters and contrast elevation from each fitted curve obtained under low and high perceptual load were compared. Paired t-tests were used to reveal significant parameter differences between the two attention conditions.

2.2.2 Results

Visual search. Visual search reaction time (*M* = 837ms, *SEM* = 33) was significantly increased and accuracy (*M* = 83%, *SEM* = 2.1) significantly reduced under high load compared to reaction time (*M* = 692ms, *SEM* = 35) and accuracy (*M* = 93%, *SEM* = .84) under low load; *t*(11) = 8.30, *p* < .001 and *t*(11) = 7.22, *p* < .001 for reaction time and accuracy scores, respectively; confirming the efficacy of the attentional manipulation.

Orientation detection. In the orientation detection task, contrast thresholds steadily increased (indicating reduced sensitivity) when the difference between the target and mean noise orientation became smaller (Figure 2.3a). This confirms the efficacy of the noise-masking procedure.

Figure 2.3

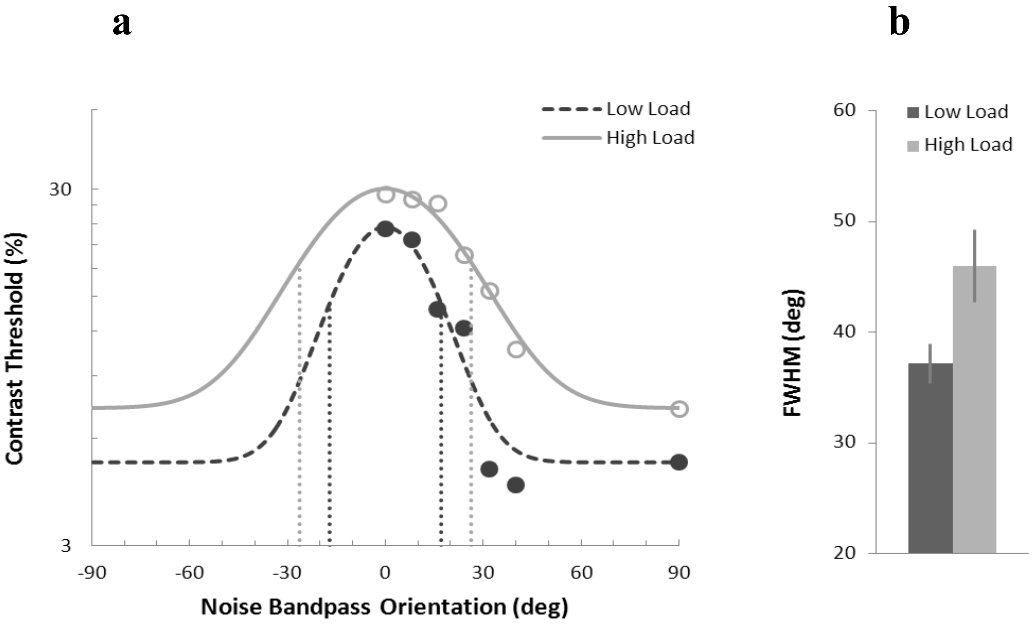

Figure 2.3. The effect of perceptual load on orientation tuning. (a) Example orientation tuning curves from one observer under low and high perceptual load in the letter search task in

Experiment 1. Contrast thresholds in the orientation detection task were fit with a modified Gaussian function. The vertical dotted lines delineate the full width of the curve at half maximum (FWHM). **(b)** Mean bandwidth across 12 observers (FWHM) in degrees. Error bars correspond to ±1 *SEM*.

The orientation tuning curves showed a robust increase in contrast threshold across all noise orientations with high compared to low load (Fig. 2.3a) and the averaged thresholds showed a significant increase from low (*M* = 10.4%, *SEM* = .90) to high load (*M* = 15.0%, *SEM* = 1.9), *t*(11) = 4.02, *p* < .01 (all t-tests in this Chapter are paired-samples t-tests).

Importantly, the threshold elevation under high load was not uniform across the different noise orientations, indicating a change in the tuning bandwidth (Figure 2.3a and b). A paired t-test on the measured bandwidth parameters (FWHM) confirmed that high perceptual load significantly increased the tuning bandwidth by 8.9° on average, *t*(11) = 3.48, *p* < .01.

2.2.3 Discussion

First of all, the results of Experiment 1 demonstrate a reduction in orientation signal strength under high perceptual load in an unrelated task. This finding is in line with previous reports of reduced perceptual sensitivity (e.g. Bahrami et al., 2008; Carmel et al., 2011; Macdonald & Lavie, 2008), as well as, reduced neural responses (e.g. Muggleton et al., 2008; Rees et al., 1997; Schwartz et al., 2000) under high perceptual load in an unrelated task.

However, in addition to the observed reduction in signal gain under high load, the psychophysical tuning curves broadened compared to low load. This suggests a novel mechanism underlying the effects of perceptual load in orientation perception which has not been discovered previously due to the limited spatial resolution of neuroimaging techniques.

2.3 Experiment 2

The results of the first experiment suggest that reduced perception or *inattentional blindness* under high load is due to both, reduced gain and broadening of the tuning. However, the manipulation of perceptual load through letter similarity in the letter search task may have involved added orientation content, and thus additional external noise content, in high compared to low load processing (since the target search among the more similar heterogeneous non-target letters required discrimination of more orientations than the less similar and homogenous non-target search in the low load condition). The added noise in the orientation detectors could directly affect orientation tuning (see Dosher & Lu, 2000a; Dosher & Lu, 2000b; Lu & Dosher, 1998). To test whether the effects of perceptual load on tuning can also be found in a load manipulation that neither varies the level of orientation information content nor the level of noise relevant for orientation perception, in Experiment 2 the letter search task was replaced with a colour-conjunction disc search task. Observers searched for a uniquely coloured target among homogeneous (low load) or heterogeneous coloured discs (high load; Fig. 2.1b), equalizing the orientation content in the low and high load conditions.

2.3.1 Method

Observers. Ten observers (8 females, 20-29 years old, *Median* = 22) participated in the second experiment. All had normal or corrected to normal vision, were naïve to the purpose of the experiment and none had participated in Experiment 1. All were recruited through the University College London subject pool.

Stimuli & Procedure. The procedure and stimuli were identical to Experiment 1, except that observers performed a colour conjunction task instead of the letter search task (see Figure 2.1b): Four discs subtending 2° of visual angle each were presented surrounding the Gabor patch and noise-mask at four fixed locations equidistant from the centre. The discs consisted of three concentric circles which were coloured red, blue or green. All three colours always appeared in a single disc but colours applied to specific circles varied, creating a total of 6 possible colour combinations. In both load conditions, participants searched the display for one of two predefined targets with a unique colour combination. One of the two targets was presented in a random half of the trials while the other target appeared in the remaining trials. Three distractors were presented in the other locations. In the low load condition the distractors were identical within each individual trial, while in the high load condition each distractor had a unique colour combination. The distractors in both conditions were randomly selected from the remaining non-target discs on each trial (with the constraint that in the high load condition distractors could not be identical).

2.3.2 Results

Visual search. The results showed that high load significantly reduced search accuracy (low load *M* = 94%, *SEM* = 1.0, high load *M* = 81%, *SEM* = 2.6), *t*(9) = 6.89, *p* < .001 and increased reaction time (low load *M* = 836ms, *SEM* = 24, high load *M* = 990ms, *SEM* = 27), *t*(9) = 8.79, *p* < .001, demonstrating that this type of attentional manipulation was effective (and produced comparable levels of performance to the letter search task used previously).

Orientation detection. As in Experiment 1, contrast thresholds were significantly increased under high perceptual load (averaged thresholds for high and low load were 15.2%, *SEM* = 1.8, and 11.2%, *SEM* = 1.1, respectively; *t*(9) = 4.14, *p* < .01). More importantly, as illustrated in Figure 2.4, under high perceptual load in the search task the tuning bandwidth (FWHM) in the orientation discrimination task was again significantly increased (13.9° on average, *t*(9) = 3.15, *p* < .05); thus, replicating the effects of perceptual load on gain and tuning with a load manipulation that involved no added orientation content.

Figure 2.4

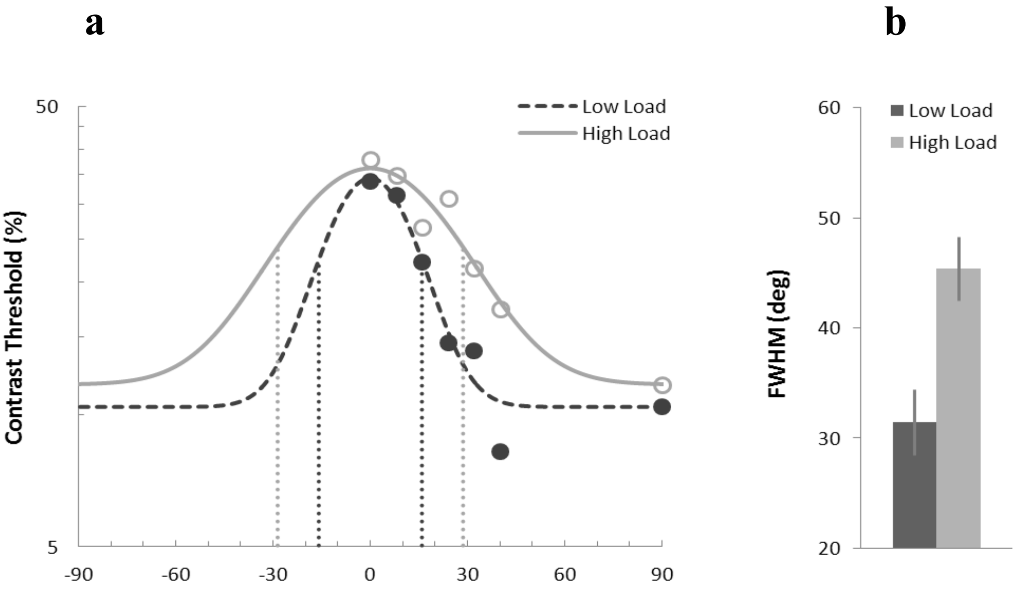

Figure 2.4. The effect of perceptual load on orientation tuning when orientation content in the search task was constant. (a) Example orientation tuning curves from one observer under low and high perceptual load in the colour-conjunction search task in Experiment 2.

The vertical, dotted lines delineate the full width of the curve at half maximum (FWHM). **(b)** Mean bandwidth (FWHM) parameters obtained under low and high load across all observers. Error bars correspond to ±1 *SEM*.

2.3.3 Discussion

The results of Experiment 2 replicate the initial finding that high perceptual load in a separate task reduces the gain and broadens the tuning of the orientation tuning profile. Furthermore, the results confirm that the effect of load on orientation responses does not depend on the orientation content in the visual search task. As performance in the letter search task of Experiment 1 was comparable to performance in the colour-based search task in Experiment 2, it can be concluded that the level of load in the task is the cause for broadened orientation tuning, irrespective of the type of task performed.

2.4 Experiment 3

Previous research has demonstrated that contrast sensitivity is not only increased at attended locations but also appears reduced at unattended locations (e.g. Pestilli & Carrasco, 2005). Experiment 3 investigated whether the effect of perceptual load on tuning can be attributed to a reduction in the perceived contrast of the noise mask. If this were the case, simply reducing the mask contrast should produce similar results to those obtained in the high load condition. To this end, observers performed the orientation discrimination task alone, with two different contrast levels for the noise mask.

2.4.1 Method

Observers. Eight observers participated in Experiment 3 (5 females, 21-29 years old, *Median* = 26), four of which had also participated in Experiment 2. All observers had normal or corrected to normal vision and were naïve to the purpose of the experiment.

Stimuli & Procedure. The stimuli and procedure were identical to those used in Experiment 1 except that observers now performed the orientation discrimination task in isolation (without any additional load task) and the mask contrast was either the same as before (15% rms contrast) or half the level (7.5% rms contrast). Mask contrast levels were blocked and their order counterbalanced across participants. Four thresholds were obtained for each noise mask contrast and noise orientation.

2.4.2 Results

Orientation detection. As expected, the amplitude of the tuning curve increased significantly with higher mask contrast (threshold under high mask contrast $M = 11\%$, $SEM = .71$, low mask contrast $M = 5\%$, $SEM = .31$; $t(7) = 12.07$, $p < .001$). However, tuning bandwidth (FWHM) remained unchanged (with high mask contrast $M = 30°$, $SEM = 3.1$, with low mask contrast $M = 34°$, $SEM = 3.5$; $t(7) = .83$, $p = .44$). Thus, tuning bandwidths for both low and high mask contrasts were comparable to those obtained under low perceptual load in Experiments 1 and 2. This result rules out a change in the effective contrast as a possible explanation for the change in tuning under high load.

Table 2.1. Summary of results from Experiments 1-3. Mean bandwidth and amplitude estimates of orientation tuning curves obtained under low and high load (Experiments 1 and 2) and low and high mask contrast (Experiment 3). High perceptual load consistently increased tuning bandwidth and amplitude parameters in Experiments 1 and 2. Increased bandwidth was not the result of an effective reduction in mask contrast under high compared to low load (Experiment 3).

	Bandwidth (FWHM) degrees			**Amplitude** % contrast		
	Low Load	**High Load**	***p***	**Low Load**	**High Load**	***p***
Experiment 1 - Letter Search	37.15 (1.81)	46.00 (3.26)	**<.01**	15.22 (1.01)	20.16 (1.71)	**<.05**
Experiment 2 - Color Search	31.42 (2.98)	45.36 (2.90)	**<.01**	16.04 (1.76)	21.43 (2.50)	**=.06**
	15% Mask	**7.5% Mask**	***p***	**15% Mask**	**7.5% Mask**	***p***
Experiment 3 - No Search Task	29.98 (3.14)	34.48 (3.50)	**=.44**	11.28 (0.71)	4.82 (0.31)	**<.001**

Numbers in parentheses indicate ±1 *SEM*; *p*-values are from paired-samples *t*-tests.

2.4.3 Discussion

Some studies have found that spatial attention does not only enhance contrast sensitivity at attended locations but also reduces the apparent contrast at unattended locations (e.g. Pestilli & Carrasco, 2005). This finding is in line with the results of Experiments 1 and 2 which showed overall increased contrast thresholds under high compared to low load. As the broadened tuning under high load was always observed in combination with the increase in contrast threshold it is possible that the reduction in contrast sensitivity may have been the determining factor of tuning width.

Experiment 3 disentangled these factors and showed that although contrast thresholds were increased when mask contrast was reduced (emulating the effective reduction in contrast accompanying high load) this did not lead to a change in tuning bandwidth. Thus, a reduction in apparent contrast under high load (or a reduction in gain) is not sufficient to produce broadened tuning.

2.5 Chapter Conclusions

Overall the present findings establish perceptual load as an important determinant of both neural sensory gain and tuning. These combined effects suggest a compelling explanation for the robust modulations of neural activity: high perceptual load has been shown to virtually eliminate neural response to unattended information over a wide range of tasks, stimuli and cortical areas (see Lavie, 2005; 2010 for review). The present results provide a powerful likely mechanism that can explain the elimination of neural signals to unattended information seen in previous studies. With both reduced signal and increased noise due to imprecise tuning in conditions of inattention under load, visual cortex response can no longer be discriminated from baseline levels of activity.

At the cellular level, neural response to orientation is determined not only by which orientation columns are stimulated but also by the strength of lateral interactions between the cell columns (Ferster & Miller, 2000). Although higher perceptual load is unlikely to modulate the response bandwidth of individual neurons (this appears to be determined simply by anatomical position; see also McAdams & Maunsell, 1999), high perceptual load could reduce the gain of individual neurons, thus leading to a modulation of lateral inhibitory interactions. Reduced lateral inhibition under high load would result in broadening of the population tuning due to reduced inhibitory inputs from neurons tuned to other orientations. Enhanced responses to normally suppressed orientations would broaden the overall population response, effectively increasing the level of noise in the sensory representation of a given orientation. Note, however, that such reduced inhibition would not necessarily result in a net increase in neural activity. In fact a general reduction in inhibition fits well with the normally observed reduction in BOLD response under high perceptual load (e.g. Rees et al., 1997; Schwartz et al., 2005). For example, Kinoshita and colleagues (Kinoshita, Gilbert & Das, 2009) demonstrated, by using optical imaging techniques that local interactions among oriented lines in early visual cortex can result in reduced inhibition. The authors showed that the previously observed facilitation in neural firing rate by collinear flankers (Kapadia et al., 1995) goes along with reduced inhibition and overall reduced local metabolic demand. Thus, a general reduction in BOLD response under high perceptual load may be the result of either, reduced inhibition (resulting in broader orientation tuning at the level of the population) or reduced excitation (resulting in reduced signal gain), or a combination of the two.

In addition, the taking up of neural capacity in conditions of high perceptual load may lead to a reduced number of neurons responding to the orientation, and thus also reduce the amount of noise cancelation that is achieved through population averaging. With a smaller population of neurons responding, the inherent variability of single neuron responses is not averaged out as well, thus resulting in a noisier response showing as broadened tuning (Gilbert & Li, 2013).

Previous studies of the effects of attention on orientation perception have typically used spatial cueing to vary the allocation of attention. This research indicated a clear effect of spatial cuing on gain (Baldassi & Verghese, 2005; Carrasco et al., 2000; Eckstein et al., 2002; McAdams & Maunsell, 1999) but has typically failed to show any effects on tuning, even when using similar noise masking techniques as that reported here (e.g. Baldassi & Verghese, 2005; Eckstein et al., 2002; Murray, Sekuler

& Bennett, 2003). The findings reported here raise the possibility that spatial cuing effects on tuning may require a higher level of perceptual load.

Indeed, this conclusion is consistent with Lu and Dosher's (1998; Dosher & Lu, 2000a, Dosher & Lu, 2000b) suggestion that the effects of spatial cuing depend on a noise exclusion mechanism with a limited capacity. However, the previous support for this suggestion may not necessarily indicate broadening of orientation tuning with higher perceptual load because the effects on noise exclusion were shown only for orientation discrimination from among white noise, and thus cannot inform about the overall orientation tuning profile, unlike the present study procedure. Moreover, the results have generalized the effects of perceptual load across a colour-based manipulation that does not vary at all the level of external noise directly involved in orientation perception (Experiment 2). The finding of broadened tuning under higher perceptual load is therefore not due to an increase in the level of external noise presented under higher load.

Another previous report indicating that orientation detection is less precise in dual versus single task conditions (Lee et al., 1999) is also consistent with the present findings. However, in Lee et al.'s study, attention had to be divided between two sets of orientations in the dual task but not in the single task condition. The single and dual-task comparison therefore did not only involve increased attentional load but also increased orientation noise, as well as added demands on memory and response.

The current manipulation of the level of perceptual load in dual task conditions only, together with the generalization of results across manipulations of perceptual load that involved no change in the level of orientation noise, avoids these pitfalls and clearly demonstrates that it is the level of perceptual load per se that is critical for the effect of attention on tuning. Moreover, the results demonstrate that engaging attention in a high perceptual load task leads to changes in gain and tuning without relying on spatial cueing or target location uncertainty. Lastly, an account in terms of feature-based attention (e.g. Ling et al., 2009; Martinez-Trujillo & Treue, 2004) is not viable either since (1) the manipulation of perceptual load did not involve any cuing of the orientation feature, (2) the orientation of the target was always constant and (3) if feature based attention was influenced by increased load in the letter-search task in Experiment 1, high load in the task should lead to increased feature based attention for the target (opposite to our results) as orientation discrimination for the target letter becomes more demanding.

In conclusion, the results from Chapter 1 demonstrate that high perceptual load in a letter or colour based visual search task leads to both reduced signal gain and broadened tuning for orientation detection. These results suggest that a combined mechanism, involving both reduced neural response gain and reduced precision of tuning, underlies the reduced sensory processing and accompanying perceptual consequences for vision (such as *inattentional blindness*) under high perceptual load.

Chapter 3

The effect of perceptual load on contextual integration: The direct tilt-illusion

3.1 Chapter Introduction

The results of the previous chapter suggest that the availability of perceptual resources plays a crucial role in determining not only the strength but also the quality of perceptual representations. The observed change in the precision of orientation responses is likely to result from modulation of the global tuning profile representing the attended orientation target in visual cortex. Martinez-Trujillo and Treue (2004), for example, showed how modulations in gain of individual neurons, without any changes in their tuning bandwidth, can result in overall sharpening of the population response. Furthermore, many neurophysiological studies have demonstrated that the tuning bandwidth of individual neurons is unaffected by changes in contrast (e.g. Troyer, Krukowski, Priebe & Miller, 1998) or attention (e.g. Martinez-Trujillo & Treue, 2004; McAdams & Maunsell, 1999; Treue & Martinez-Trujillo, 1999). Thus, it appears that any change in the precision of orientation representations must occur through relatively large scale, lateral interactions among orientation columns, modulating the gain of individual neurons preferring a given orientation. This same mechanism is thought to underlie contextual interactions that alter orientation perception (e.g. Blakemore & Tobin, 1972). Supporting evidence has been provided by a recent study using TMS to assess mechanisms underlying orientation processing in humans (Ling, Pearson & Blake, 2009). The authors showed that the bandwidth of psychophysical orientation tuning curves was unaffected by applying TMS to early visual areas. However, contextual effects of overlapping orientations (in this case for the direct tilt-illusion) were significantly reduced by TMS to the same cortical location.

Together, these findings provide a clear hypothesis for the effects of perceptual load on orientation processing. If high perceptual load in an unrelated task indeed reduces the precision of orientation responses at a global level, it should also affect contextual integration of orientations. More specifically, with limited perceptual resources under high load, contextual effects should be reduced if high load reduces lateral inhibition. For the case of the tilt-illusion, for example, the strength of the illusion should be reduced under high perceptual load, leading to a more veridical percept.

Recently, it has been demonstrated that a reliable tilt-illusion can be observed, even when observers are completely unaware of the orientation of the context (Clifford & Harris, 2005). Here the location of a tilted inducer surrounding a central target was masked by a high contrast backward mask, precluding conscious identification of the inducer orientation. This finding provides a useful methodology to reduce perceptual demands associated with distinguishing target from inducer when they are presented close together (or attention is focused elsewhere). Furthermore, it can reveal load effects on the interaction between unconscious visual context and the consciously perceived target.

3.2 Experiment 4

The purpose of Experiment 4 was to establish whether perception of the direct tilt-illusion is modulated by the level of perceptual load in an unrelated task.

The findings from the preceding empirical chapter suggest that the effects of load are at least partly mediated by a reduction in the precision of orientation specific responses (characterized by broadened orientation tuning at the level of the neural population).

As detailed in the introduction above, the orientation tuning properties of single neurons appear unaffected by changes in contrast (e.g. Troyer, Krukowski, Priebe & Miller, 1998), attentional manipulations (e.g. McAdams & Maunsell, 1999) or suppression by TMS applied to early visual cortex (Ling, Pearson & Blake, 2009). In contrast, other aspects of orientation processing, such as contextual effects (and the tilt-illusion described here) do depend on large scale interactions at the level of the neural population representing a given stimulus and are, for example, influenced by TMS (Ling, Pearson & Blake, 2009). The perceived tilt-repulsion has been attributed to intracortical inhibition among neurons selective to orientation (Carpenter & Blakemore, 1973). Therefore, if perceptual load indeed influences the global balance of inhibition between orientation detectors in early visual cortex, this should modulate the strength of the perceived tilt-illusion. This hypothesis leads to the counterintuitive prediction that for the case of contextual illusions involving orientation processing, withdrawing attentional resources with a high perceptual load task should lead to a more veridical percept (a reduction in the magnitude of the illusion).

In order to be able to relate the findings to those of the previous experiments, the design of Experiment 4 was similar to that of Experiment 1. Firstly, the same letter search task was used to modulate perceptual load and, secondly, as in Experiments 1 to 3, a band-pass filtered noise mask was used. The noise mask was, again, superimposed onto a vertical target stimulus at fixation. However, in this case the orientations in the noise mask were always tilted either to the left or the right from vertical in order to induce the tilt-illusion. Furthermore, the contrast of the noise mask was varied (similar to Experiment 3) to evaluate the efficacy of the mask in inducing the tilt-illusion (as increased inducer contrast has been shown to increase the tilt-illusion, e.g. Blake et al., 1985) and to assess whether any effect of perceptual load on the strength of the illusion could be accounted for by an apparent change in stimulus contrast under high load.

3.2.1 Method

Observers. Twelve observers, (8 females, 19 to 24 years of age, *Median* = 22.5) participated in Experiment 4. All had normal or corrected to normal vision and were recruited through the University College London subject pool. All observers were naïve to the purpose of the experiment.

Apparatus & Stimuli. Stimuli for all experiments in this chapter were created using MatLab (v.R2010a, The MathWorks, Nattick, MA) and Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) and presented on a 21.5” monitor (1920 x 1080 pixel resolution, 60 Hz refresh rate) in a dimly lit room. Viewing distance was fixed at 57cm using a chin rest. Stimuli were presented on a mid-grey background at the centre of the screen. Observers searched for a target letter (X or N) among seven Os (in the low load condition) or among seven unique angular letters, randomly selected from K, E, F, H, L, M, T, W, Y or Z (in the high load condition) while simultaneously judging the orientation (clockwise vs. counter-clockwise from vertical) of a Gabor patch presented at fixation (Figure 3.1). The Gabor patch (4° in diameter, 1.5 cpd, 25% Michelson contrast) was embedded within a noise patch (4° in diameter) which was band-pass filtered in the orientation (10° bandwidth) and spatial frequency (4-10 cpd bandwidth) domains with a mean spatial frequency of 7 cpd and acted as an inducer for the tilt-illusion. The mean orientation of the noise was always centred on 35° clockwise (as in Figure 3.1) or 35° counter-clockwise from vertical and its contrast was either 9% or 18% (rms contrast; Figure 3.1c). The 35° angle was chosen as this orientation difference was found to induce maximal tilt-repulsion when target and inducer are superimposed (Blake, Holopigian & Jauch, 1985).

Figure 3.1

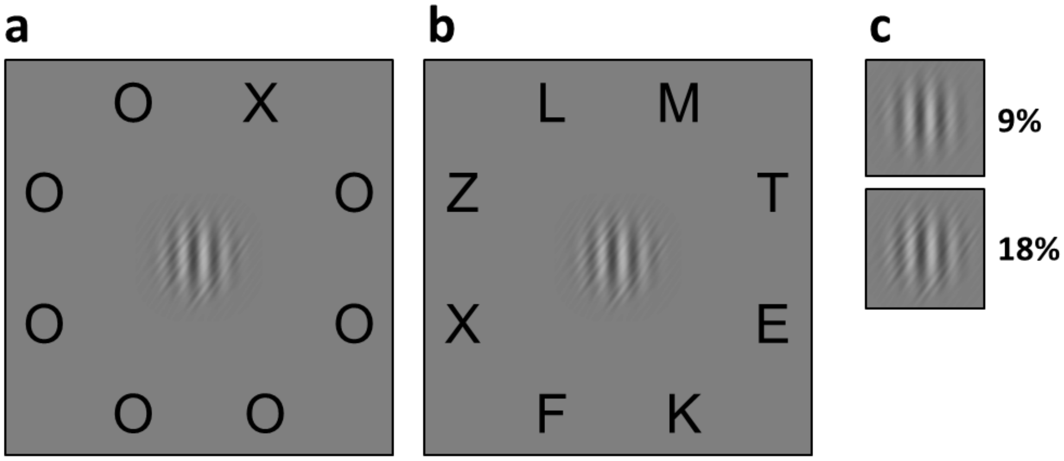

Figure 3.1. Schematic illustration of stimulus displays in Experiment 4. Observers performed a letter search task under **(a)** low and **(b)** high perceptual load while judging the orientation of a Gabor patch at fixation, here shown at 0° (vertical), in order to estimate the perceived subjective vertical. The Gabor patch was embedded in a noise mask which was tilted either 35° clockwise or 35° counter-clockwise from vertical in order to induce a tilt-illusion. **(c)** The contrast of the noise inducer was either 9% or 18% (rms contrast). The illusion increases with inducer contrast (illustrated above, at 18% the vertical target looks slightly more tilted to the left than at 9% contrast).

Procedure. Before the experiment all observers were shown still frames of the stimuli and explicitly instructed to only respond to the target letter and the orientation of the target Gabor while ignoring the orientation of the inducer (which was described as “a mesh of fine lines sometimes visible in the background”). Observers were instructed to prioritize the letter search task over the orientation judgment and were asked to respond as quickly and accurately as possible.

Each trial began with presentation of a white fixation point at the centre of the screen for 700ms. After a delay of 500ms, all stimuli were presented simultaneously for 167ms. Following stimulus presentation, a blank screen prompted observers to respond by first pressing one of two designated keys on the left side of the keyboard, indicating the target letter (X or N) and then use the left and right arrow keys on the right side of the keyboard to indicate the perceived orientation of the Gabor in the centre.

All observers completed eight blocks, four with the high load letter search and four with the low load letter search. In each block, two randomly interleaved staircases (QUEST; Watson & Pelli, 1983) were used to vary the orientation of the target Gabor in order to determine the orientation perceived as vertical when the inducer was either tilted 35° clockwise or 35° counter-clockwise from vertical. Thus, the inducer and target orientations were unpredictable on each trial, precluding systematic response biases or adaptation effects. Within a block, for one of the staircases the contrast of the inducer was always 9%, for the other it was always 18% (rms contrast). The order of staircases (inducer contrast and inducer orientation) was counterbalanced across blocks within a load condition and the order of high and low load blocks was counterbalanced across observers.

3.2.2 Results

Visual search. Performance in the letter search task was significantly reduced under high perceptual load. Observers responded less accurately under high (*M* = 82%, *SEM* = 2.02) compared to low load (*M* = 97%, *SEM* = .84; *t*(11) = 9.48, $p < .001$) and reaction times were significantly longer under high (*M* = 1052ms, *SEM* = 35) compared to low load (*M* = 789ms, *SEM* = 43; *t*(11) = 18.24, $p < .001$). These results confirm the efficacy of the attentional manipulation.

Orientation judgment. Average tilt-repulsion biases were calculated separately for each load condition and inducer contrast. Biases from both inducer orientations (35° clockwise and 35° counter-clockwise) were averaged in order to eliminate any constant errors (e.g. perceiving a vertical orientation without any inducer as tilted):

$$Bias^\circ = \frac{\mathrm{PSE}_{\mathrm{pos}} - \mathrm{PSE}_{\mathrm{neg}}}{2} \qquad \text{Equation 1}$$

Here, PSE is the point-of-subject-equality where the target appears vertical in the presence of the positive ($_{\mathrm{pos}}$) or negative ($_{\mathrm{neg}}$) inducer orientation (+35° or -35° from vertical). A two factor, within-subject repeated measures ANOVA with the factors Load and Inducer contrast revealed a significant main effect of Load, *F*(1,11) = 5.52, $p < .05$, $\eta_p^2 = .33$, indicating that mean biases were significantly smaller under high perceptual load (*M* = 1.95°, *SEM* = .36) compared to low load (*M* = 2.71°, *SEM* = .32), as well as a main effect of Inducer contrast, *F*(1,11) = 94.03, $p < .001$, $\eta_p^2 = .90$, demonstrating that biases were significantly greater at the higher, 18% inducer contrast (*M* = 3.49°, *SEM* = .29) compared to the lower, 9% inducer contrast (*M* =

1.17°, *SEM* = .21) (Figure 3.2). The interaction between inducer contrast and load was not significant, *F*(1,11) = .68, *p* = .43, indicating that the reduction in bias under high load did not depend on inducer contrast.

Figure 3.2

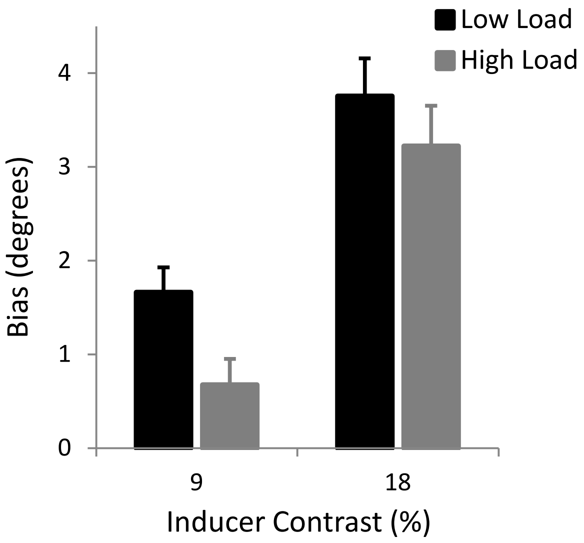

Figure 3.2. Tilt-repulsion bias as a function of inducer contrast and level of load in the letter search task in Experiment 4. Error bars indicate 1 standard error of the mean (*SEM).*

Firstly, the results show that increasing inducer contrast increased the strength of the tilt-illusion, confirming the efficacy of the inducer and target stimulus parameters in producing the illusion, regardless of the attentional modulation. Moreover, the magnitude of the overall increase in bias (of roughly 2°) with contrast is comparable to a previous study using the same type of target and inducer (Blake et al., 1985). Secondly, the tilt-illusion was significantly reduced under high perceptual load in the letter-search task. Lastly, the effect of high load is unlikely to result from a reduction in effective contrast of the inducer as the effect appears larger at the low (9%) compared to the high (18%) mask contrast (see Figure 3.2).

3.2.3 Discussion

The results indicate that the magnitude of the perceived tilt-repulsion depends on the level of perceptual load in an unrelated task. The finding of reduced tilt-repulsion under high perceptual load supports the hypothesis that broadened orientation tuning under high load (observed in Chapter 2) may be the result of reduced inhibitory interactions in early visual cortex. Weakened inhibition among lateral connections under high load should reduce the magnitude of the perceived tilt-illusion. However, a change in lateral inhibition could also result in reduced visual resolution for unattended stimuli under high load, in line with previous findings demonstrating increased resolution at attended compared to unattended locations (Yeshurun & Carrasco, 1998). This finding has been interpreted in terms of a change in the size of spatial filters (or receptive fields) at the attended location. In the current experiment, an increase in the size of spatial filters at fixation under high load would primarily reduce the effect of the high spatial frequency inducer (see Figure 3.1) and thus, reduce the magnitude of the tilt-illusion under high load.

Both of these accounts would indicate that under the right circumstances a reduction in attentional resources can indeed result in a more veridical percept.

Alternatively, the observed effect may be a result of the reduced ability to discriminate orientations under high load. As in the current experiment target and inducer were superimposed, observers had to discriminate target from inducer first in order to make a judgment with regard to the orientation of the target. The results from the previous chapter demonstrated that the ability to distinguish the target from orientation noise is generally reduced under high perceptual load. Therefore, it is possible that on some trials observers erroneously responded to the inducer orientation instead of the target orientation. Any response to the inducer instead of the target would result in an apparent reduction in overall tilt repulsion. Therefore, the strength of the illusion would appear to be reduced under high perceptual load due to less precise measurement, but not due to a change in the perceived illusion itself. The next experiment, Experiment 5, addressed this concern by reducing inducer visibility and spatially separating the inducer from the target.

3.3 Experiment 5

The purpose of Experiment 5 was two-fold. Firstly, it was designed to overcome the limitations of the previous experiment by using a paradigm which eliminated the effect of reduced discriminability (of target and inducer orientations) on the tilt-illusion. In this way the direct effect of load on the perceived tilt-repulsion bias could be investigated, excluding any influence of target discriminability on responses (i.e. responding to the inducer orientation instead of the target) or on perception of the target (i.e. additional noise from the inducer in the same location, or different effects on target and inducer due to changes in the size of perceptual filters (e.g. Yeshurun & Carrasco, 1998)). To achieve this goal, the contrast of the central target grating was increased while the inducer was no longer overlapping with the target grating but was now surrounding it. Furthermore, target and inducer had the same spatial frequency. Spatial separation of inducer and target should aid observers in identifying the target at the centre without having to discriminate target from inducer before judging the target’s orientation.

Secondly, the experiment investigated whether perceptual load would influence the perceived tilt-illusion even when observers were not consciously aware of the inducer (e.g. Clifford & Harris, 2005; see Chapter Introduction). To this end, and to completely eliminate any demands on discriminating target and inducer orientations, a high contrast, random noise, backward-mask was used in the location of the inducer. The mask appeared immediately after presenting target and inducer together, rendering the inducer "invisible". As observers were unaware of the orientation of the inducer, this procedure also eliminated any response biases that may have existed for a given inducer orientation in the previous experiment (e.g. some of the observers may have been familiar with the tilt-illusion phenomenon and may have tried to counteract its effects).

3.3.1 Method

Observers. Ten observers participated in Experiment 5 (6 females, 18 to 26 years old, *Median* = 21.5). All had normal or corrected to normal vision, were naïve as to the purpose of the experiment and recruited through the University College London subject pool.

Stimuli. Stimuli were presented on a medium-grey background in a dimly lit room (Figure 3.3). Viewing distance was fixed at 57cm using a chin rest. Observers performed a dual-task, judging the orientation of a grating at fixation (clockwise or counter-clockwise from vertical), while simultaneously performing a letter search task. The letters (black, 0.9° in width and 1.1° in height) were equally spaced at 3.4° from the centre. In the low load blocks observers searched for an X or N among seven Os, while in the high load blocks they searched for an X or N among seven different angular letters (randomly selected from K, E, F, H, L, M, T, W, Y or Z). The central grating (2.2° in diameter, 5 cpd, 72% Michelson contrast) was surrounded by a grating annulus (5.2° in diameter, 1.5° width, 5 cpd, 50% Michelson contrast) which was either oriented 35° clockwise or 35° counter-clockwise from vertical. A random dot noise mask (same dimensions as annulus, pixel values randomly selected from a standard normal distribution) was used to mask the orientation of the annulus.

Figure 3.3

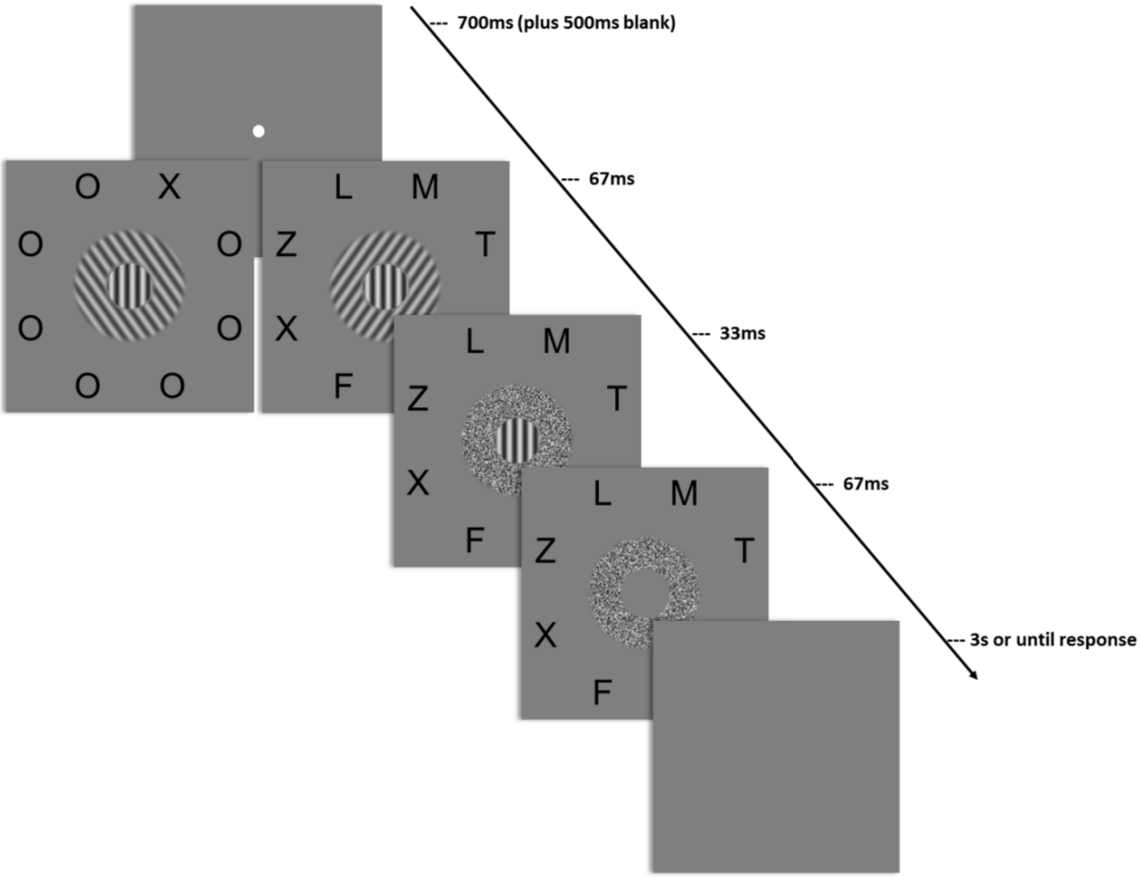

Figure 3.3. Schematic illustration of stimuli and trial sequence in Experiment 5. Observers searched for a target letter among Os (low load) or angular letters (high load) while judging the perceived orientation of the grating at the centre. The annulus surrounding the central grating contained a lower contrast grating which was always tilted 35° clockwise or

35° counter-clockwise from vertical (randomised over trials). The tilt of the surround causes the inner grating to appear tilted in the opposite direction (tilt-repulsion) even when the surround is rendered invisible by the subsequent high contrast noise mask.

Procedure.

Baseline bias estimation. Before the main experiment, all observers completed two experimental blocks to estimate the baseline tilt-illusion bias without an additional load task, in order to determine the efficacy of the “invisible” inducer to produce the tilt-illusion. Each trial began with presentation of a white fixation dot, for 700ms, at the centre of the screen. Following a delay of 500ms the letter search task was presented together with the central target grating and the annulus grating (Figure 3.3). The annulus grating was then replaced by a noise mask (after 67ms) while the central target grating and the letters remained on screen. After 33ms the central target was removed, leaving the letters and noise mask for another 67ms on the screen. Thus, the letter search task was presented for 167ms in total while the central target was visible for a total of 100ms. In order to simulate the stimulus conditions of the main experiment, the letter search task was always presented during baseline estimation of the tilt-illusion (high or low load, counterbalanced across observers). However, observers were explicitly instructed to completely ignore the letters and the surround annulus (i.e. the perceived noise mask), and only judge the orientation of the central grating (clockwise vs. counter-clockwise from vertical). In each block, two randomly interleaved staircases adaptively varied the orientation of the central grating in order to find the orientation perceived to be vertical (QUEST). For one of the staircases the annulus orientation was always 35° clockwise, for the other it was 35° counter-clockwise. Each staircase started at 0° (vertical). On each trial observers pressed a predefined key on the left side of the keyboard (as if responding to the target letter) before giving the response to the orientation of the target grating (using the left and right arrow keys on the right side of the keyboard). Importantly, observers at this point were unaware of the role the letters would have in the main experiment.

Main experiment. The stimuli and trial sequence of the main experiment were identical to those of the baseline estimation experiment, except that observers now performed the letter search task in addition to judging the orientation of the central grating. Moreover, observers were instructed to prioritize the letter search task over the orientation judgment. Each observer completed four blocks under low and high load in the letter search task. The same staircase procedure as in the baseline estimation was used to find the subjective vertical for each load condition and annulus orientation (+35° or -35° from vertical). Observers were instructed to respond as quickly and accurately as possible by first pressing one of two keys on the left side of the keyboard, indicating the perceived target letter (X or N), and then pressing either the left or right arrow key on the right side of the keyboard to indicate the orientation of the central grating.

Annulus visibility and control experiment. After completion of the main experiment, all observers were questioned about whether they had perceived the tilted annulus surrounding the target grating. Questions were always asked in the following order: 1. “How would you best describe what you saw in the area between the grating in the middle and the letters surrounding it?” 2. “Did you notice any bars or lines in this area, similar to those in the centre?” 3. “Did you notice any orientation or direction in this area?”

Observers were then shown a still frame of the annulus grating together with the central grating and the letters and were told that the tilted annulus was present on each trial but had been rendered “invisible” by the subsequent noise mask. With this in mind, observers performed a control experiment where they searched for the target letter (X or N) under low and high load, as before, while at the same time judging the orientation of the annulus grating (2-AFC) in the presence of the noise mask and ignoring the orientation of the central grating. The stimuli and trial sequence were identical to those of the main experiment, except that the orientation of the central grating was randomly selected (from a range of orientations between 45° clockwise and 45° counter-clockwise from vertical) on each trial. Observers completed two blocks (one for each load level in the letter search task). As before, on a random half of the trials in each block the annulus orientation was 35° clockwise, on the other half it was 35° counter-clockwise from vertical.

3.3.2 Results

Baseline tilt-repulsion bias. Biases were taken as half of the difference of the orientation perceived as vertical when the surround was tilted clockwise, and that when the surround was tilted counter-clockwise (as in Experiment 4, see Equation 1). Without the additional letter search task, the tilted (and masked) annulus induced a consistent tilt-repulsion bias in all observers (ranging from .57° to 1.97°). The average baseline bias was 1.07° (*SEM* = .15), which was significantly different from zero; $t(9) = 7.26$, $p < .001$. This finding confirms that the annulus grating was effective in producing a reliable tilt-illusion, even when presented very briefly (for 67ms) and rendered “invisible” by a subsequent noise mask (see control experiment below).

Main experiment.

<u>Visual search.</u> Under high perceptual load in the letter search task, accuracy (*M* = 85%, *SEM* = 2.33) was significantly reduced while reaction time (*M* = 888ms, *SEM* = 54) was significantly increased compared to accuracy (*M* = 96%, *SEM* = .81) and reaction time (*M* = 734, *SEM* =54) under low load; $t(9) = 4.49$, $p < .01$ and $t(9) = 3.98$, $p < .01$, for letter search accuracy and reaction time, respectively. This confirms the effectiveness of the attentional manipulation.

Orientation judgment. For the main experiment, tilt-repulsion biases were calculated separately for each of the two load conditions. A one-way repeated measures ANOVA with Load as a factor with three levels (no load, low load and high load) revealed a significant main effect of load, $F(2,18) = 3.73$, $p < .05$, $\eta_p^2 = .29$, indicating that the bias was strongest under high load.

Figure 3.4

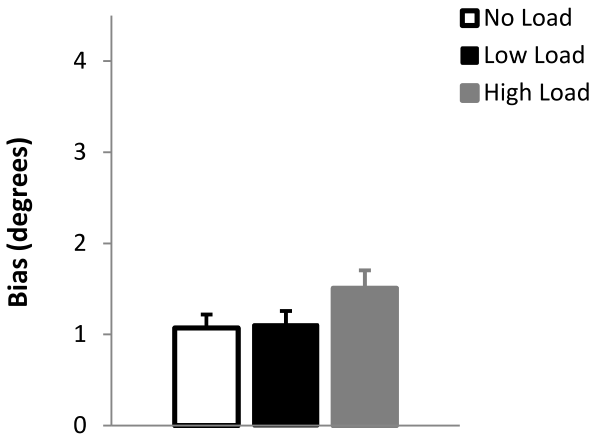

Figure 3.4. Average tilt-repulsion biases under low (black) and high load (grey) in the letter search task, and without an additional task (white) in Experiment 5. Error bars indicate 1 *SEM*.

A planned comparison between low (*M* = 1.10°, *SEM* = .16) and high load (*M* = 1.51°, *SEM* = .20) confirmed this finding, showing that the strength of the tilt-illusion significantly increased under high compared to low perceptual load; $F(1,9) =$

5.17, $p < .05$, $\eta_p^2 = .37$. Furthermore, tilt-illusion biases were not significantly different under low load compared to the baseline condition, where no additional task was performed (no load); $F(1,9) = .05$, $p = .83$ (Figure 3.4).

Visibility of annulus orientation.

Questionnaire. After completing the main experiment, all observers answered three questions to assess visibility of the surround orientation. Answers to the 1st question, “How would you best describe what you saw in the area between the grating in the middle and the letters surrounding it?”, ranged from “grey”, “random dots”, “black and white static” to “snow like” or “flower like patterns”, but none of the observers indicated anything similar to the appearance of the grating. Furthermore, none of the observers responded positively to the 2nd or 3rd question (“Did you notice any bars or lines in this area, similar to those in the centre?” and “Did you notice any orientation or direction in this area?”).

Control Experiment. In order to test if perception of the annulus orientation was strictly unconscious while performing the load task, and not simply below some arbitrary response criterion, observers were tested in a 2-AFC control experiment, judging the orientation of the annulus instead of the central grating, while performing the same letter search task as in the main experiment. In both load conditions for the letter search, observers did not perform significantly different from chance (50% correct) for judging the orientation of the annulus: Average accuracies for correct orientation judgments (left or right from vertical) were 55.7% ($SEM = 3.47$; $t(9) = 1.65$, $p = .13$) under low load and 52.8% ($SEM = 2.50$; $t(9) = 1.11$, $p = .30$) under high load. Furthermore, visibility also did not differ significantly between the two load conditions; $t(9) = 1.21$, $p = .26$, paired-samples t-test.

3.3.3 Discussion

The results of Experiment 5 indicate that high perceptual load increases the magnitude of the tilt-illusion when target and inducer do not have to be perceptually discriminated. This finding contrasts with the results of Experiment 4 which showed a reduction in the strength of the illusion under high load when target and inducer had to be discriminated. However, in Experiment 5 the tilt-illusion increased under high perceptual load even when observers were unaware of the orientation of the inducer (or its presence). Furthermore, the strength of the tilt-illusion did not differ between the baseline condition (where no additional load task was performed) and the low load condition, indicating that the additional demands associated with performing a dual task compared to a single task were insufficient to have an effect on the perceived illusion. Lastly, by using the same letter search task as in the previous experiment, the results rule out effects based on the physical layout of the letter search task which may have provided more vertical reference lines under high load (where distractors are angular letters) compared to low load (where all distractors were Os) which could have been used to better estimate vertical orientation.

Together, the findings indicate that the modulatory effect of perceptual load on the tilt illusion is not fixed and critically depends on the stimulus configuration (spatially overlapping or separated target and inducer) or the conscious versus unconscious perception of the inducer. Experiment 6 and 7 were designed to distinguish between these factors.

3.4 Experiment 6

The aim of Experiment 6 was to distinguish between the two possible factors that could have caused a reversal in the results from Experiment 4 (reduced tilt-illusion under high load) to Experiment 5 (increased tilt-illusion under high load).

The first of these factors is the spatial separation between target and inducer. In Experiment 4, target and inducer were spatially overlapping while in Experiment 5 they were separated. The second factor comprises the physical changes to the stimuli, such as contrast and noisiness of the inducer, as well as, the mask used to degrade visibility of the inducer orientation. In Experiment 4, target and inducer had been presented at low contrast (25% Michelson contrast for the target and 9% or 18% rms contrast for the inducer), the target was a regular sine wave Gabor patch and the inducer consisted of orientation band-passed noise. In Experiment 5, however, target and inducer were both regular sine wave gratings presented at high contrast (72% and 50% Michelson contrast, respectively) and, in addition, a high contrast noise mask was used to render the annulus orientation invisible.

Experiment 6 was designed to measure the influence of the first factor, spatial separation of target and inducer stimuli. To this end, all aspects of Experiment 4 were replicated with the exception of target and inducer locations which were now spatially separated. If the effect of high perceptual load in reducing the perceived tilt-illusion relies on the spatial superposition of target and inducer, the effect should be eliminated when they are spatially separated.

3.4.1 Method

Observers. Ten new observers participated in Experiment 6 (7 females, 19 to 32 years old, *Median* = 22). All had normal or corrected to normal vision, were recruited through the University College London subject pool and were naïve to the purpose of the experiment.

Stimuli & Procedure. All aspects of the stimuli and procedure were identical to those of Experiment 4, except that the noise inducer, which was superimposed on the target in Experiment 4, was now spatially separated from the target, surrounding it instead of overlapping with it (Figure 3.5b). Target and noise inducer had the same spatial frequency and orientation content as in Experiment 4. As previously, the tilt-illusion bias was measured with two inducer contrasts (9% and 18% rms contrast).

Figure 3.5

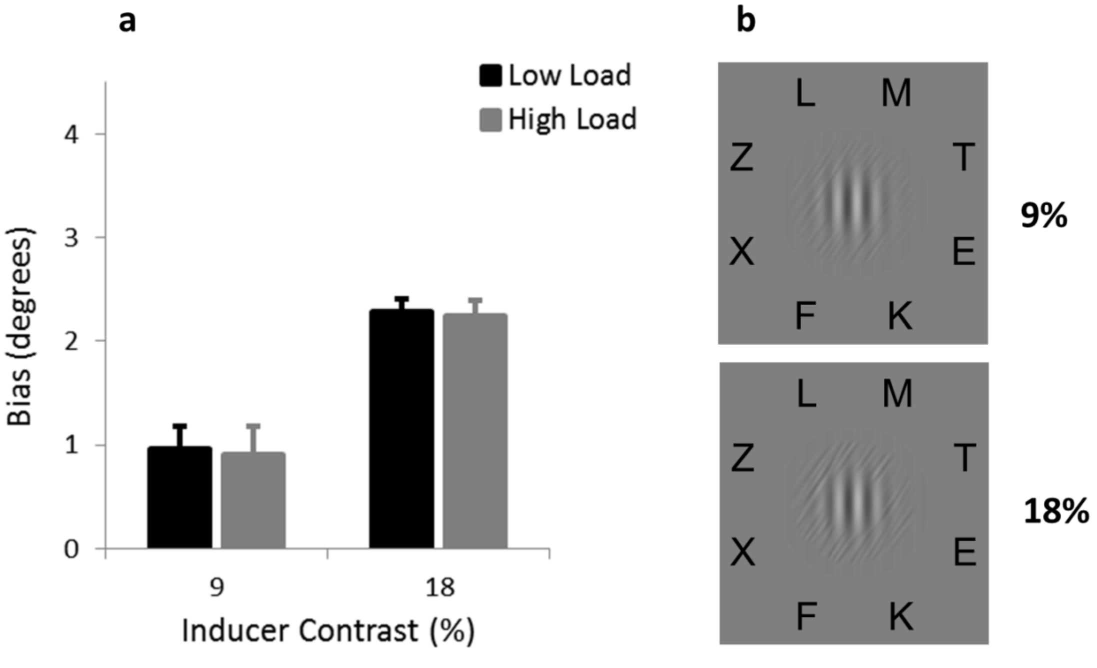

Figure 3.5. Tilt-repulsion bias and schematic illustration of stimuli in Experiment 6. (a) Tilt-repulsion bias as a function of inducer contrast and level of load in the letter search task. Error bars indicate 1 *SEM*. (**b)** Schematic stimulus display (high load condition) in

Experiment 6. The inducer (here tilted 35° clockwise) surrounded the target at the centre (here vertical). Inducer contrast (rms) was either 9% or 18%.

3.4.2 Results

Visual search. High perceptual load in the letter-search task significantly reduced accuracy (*M* = 80%, *SEM* = 3.51) and increased reaction time (*M* = 1025ms, *SEM* = 89) compared to accuracy (*M* = 94%, *SEM* = 2.81) and reaction time (*M* = 857ms, *SEM* = 74) under low load; $t(9) = 5.40$, $p < .001$ and $t(9) = 4.97$, $p < .001$ for the load effects on accuracy and reaction time, respectively.

Orientation judgment. A two-factor, within subject, repeated measures ANOVA with the factors load and inducer contrast revealed a significant main effect of inducer contrast, $F(1,9) = 70.59$, $p < .001$, $\eta_p^2 = .89$, but no effect of load and also no interaction between inducer contrast and load.

Firstly, this finding confirms that presenting the inducer in the surround still produces a reliable tilt-illusion. Secondly, the strength of the illusion was modulated by the contrast of the inducer as was the case in Experiment 4. Lastly, the level of load in the letter-search task had no effect on the strength of the perceived illusion.

3.4.3 Discussion

The results of Experiment 6 indicate that the reduction in the perceived tilt-illusion observed in Experiment 4 depends on the spatial superposition of target and inducer. This suggests that reducing the demands on target and inducer discriminability eliminates the effect of perceptual load in reducing the strength of the illusion. However, separating target and inducer did not automatically lead to an increase in the illusion under high load as was observed in Experiment 5. Therefore, in the next experiment target and inducer were again separated, however this time all other stimulus parameters were identical to those in Experiment 5 with the exception that the inducer orientation was presented consciously (no backward mask was used).

3.5 Experiment 7

As the previous experiment confirmed that spatially separating target and inducer does not automatically lead to an increase in the tilt-illusion, the purpose of Experiment 7 was to test if the increase in the perceived tilt-illusion under high perceptual load in Experiment 5 was due to the specific stimulus properties of the target and inducer (i.e. the high contrast gratings used for target and inducer in Experiment 5 compared to the low contrast target gratings and noise inducers used in Experiments 4 and 6).

3.5.1 Method

Observers. For Experiment 7, ten new observers were recruited from the University College London subject pool (19 to 35 years of age, *Median* = 24.5). All were naïve to the purpose of the experiment and had normal or corrected to normal vision.

Stimuli & Procedure. The stimuli and procedure were identical to those in Experiment 5, except that the noise mask was removed and the duration of the surround presentation was increased from 67ms to 167ms (as in Experiment 4; Figure 3.6). Observers were again instructed to ignore the oriented bars surrounding the target stimulus (the inducer) and only respond to the letter search and central target orientation, while always prioritizing the letter search as the primary task. As previously, baseline tilt-illusion biases were obtained in addition to the biases under low and high perceptual load.

Figure 3.6

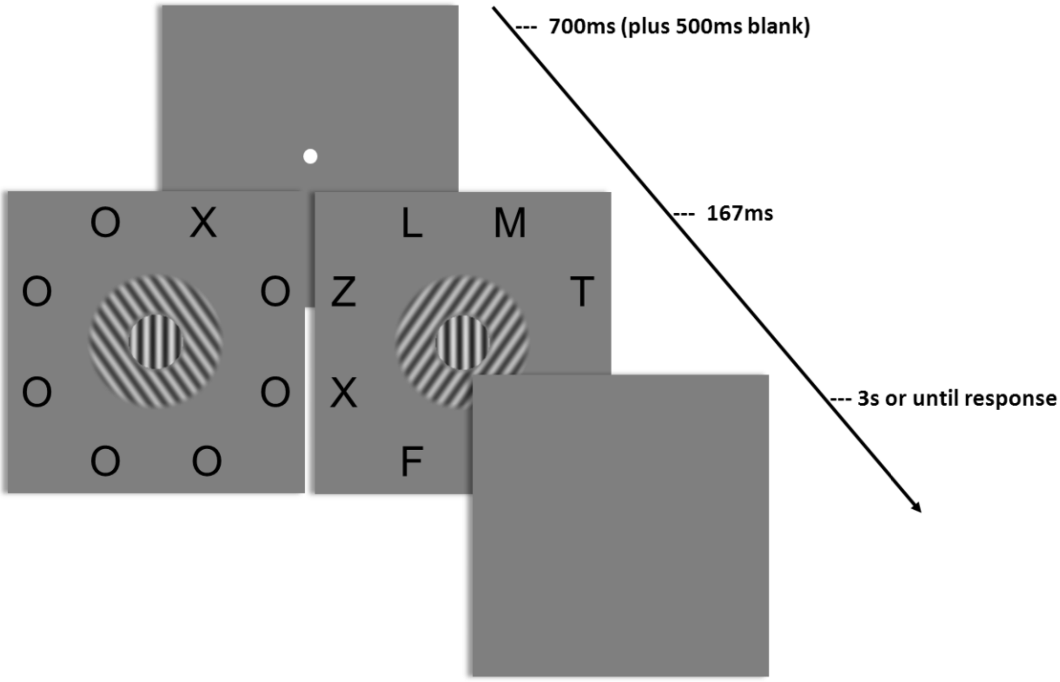

Figure 3.6. Trial sequence in Experiment 7. Observes performed the same letter search task as in the previous experiment while judging the orientation of the grating at the centre. However, no mask was presented after stimulus presentation.

3.5.2 Results

Visual search. As previously, high perceptual load in the letter-search task significantly reduced accuracy (*M* = 86%, *SEM* = 2.61) and increased reaction time (*M* = 1069ms, *SEM* = 71) compared to accuracy (*M* = 97%, *SEM* = .38) and reaction time (*M* = 923ms, *SEM* = 70) under low load; $t(9) = 3.87$, $p < .01$ and $t(9) = 4.25$, $p < .01$, for the load effects on accuracy and reaction time, respectively.

Orientation judgment. Tilt-repulsion biases were calculated for each load condition in the same way as in the three previous experiments (Figure 3.7). A one-way repeated measures ANOVA with Load as a factor with three levels (no load, low load and high load) showed no significant effect of load; *F*(2,18) = 1.29, *p* = .30. This indicates that the level of perceptual load in the letter search task had no effect on the strength of the perceived tilt-illusion.

Figure 3.7

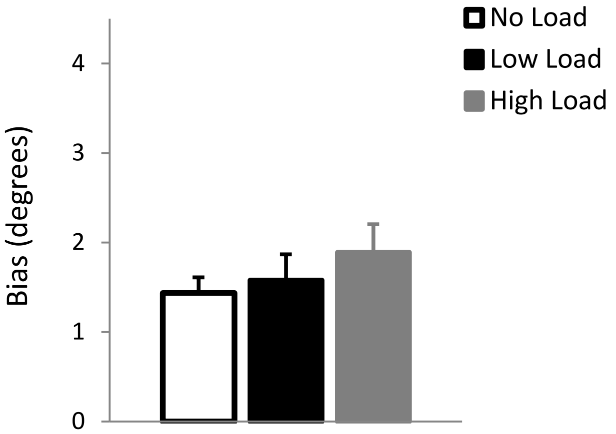

Figure 3.7. Average tilt-repulsion biases in Experiment 7, under low (black) and high load (grey) in the letter search task, and without an additional task (white). Error bars indicate 1 *SEM*.

To assess whether the use of the high contrast backward mask in Experiment 5 influenced the strength of the tilt-illusion *per se* (when no additional load task was performed), an independent-samples t-test compared the obtained baseline tilt-illusion biases from the current experiment to those in Experiment 5 (recall that Experiment 5 had identical stimulus parameters with the exception of the backward mask at the location of the inducer). Baseline biases in Experiment 5 (*M* = 1.07°, *SEM* = .15) and the current experiment (*M* = 1.44°, *SEM* = .17) did not significantly differ; *t*(18) = 1.60, *p* = .13. This suggests that the significant increase in the perceived magnitude of the tilt-illusion in Experiment 5 is due to the interaction between load and the visibility of the inducer orientation and not the physical presence of the mask alone.

3.5.3 Discussion

The results of Experiment 7 confirm that the effects of perceptual load on the magnitude of the tilt-illusion do not depend on the physical stimulus parameters of inducer and mask in terms of contrast and noisiness. In line with Experiment 6 which used the same centre-surround configuration but much reduced target and inducer contrasts, perceptual load had no effect on the tilt-illusion. Therefore, the significant increase in the strength of the illusion under high load in Experiment 5 cannot be attributed to changes in the physical stimulus parameters (between Experiments 4 and 5) but must depend on the presence of the back-ward mask, rendering the inducer orientation unconscious. As Experiment 4 showed a reduction in tilt-illusion under high load when target and inducer overlapped, it appears that the increase in the magnitude of the illusion found in Experiment 5 requires both, spatial separation of target and inducer as well as unconscious presentation of the inducer.

It was not feasible to test the remaining stimulus combination of spatially superimposing the target with an unconscious inducer, as in this case masking the inducer orientation alone, without also masking the target, would be impossible (in principle, however, this could be accomplished using binocular rivalry and continuous flash suppression (Tsuchiya & Koch, 2005), suggesting an interesting avenue for future research).

Due to the design of the experiments any use of strategy by the observers is very unlikely. By randomly interleaving the staircases measuring orientation bias, the inducer and target orientation was unpredictable on a given trial. Furthermore, observers were always instructed to ignore the inducer and prioritize the letter search task over the orientation discrimination task.

3.6 Chapter Conclusions

In summary, the level of perceptual load in a separate task can influence contextual effects arising from interactions among orientations. When perceptual load in the task is high and target and inducer are co-localized, the magnitude of the tilt-illusion is reduced. However, when target and inducer are presented in separate spatial locations, perceptual load does not influence the magnitude of the perceived illusion. Lastly, when target and inducer are spatially separated but the inducer orientation is rendered unconscious, high perceptual load increases the strength of the illusion.

This series of results suggests that high perceptual load reduces the impact of consciously perceived but task irrelevant stimulus elements. In Experiment 4 the effect of the inducer was reduced (compared to Experiment 6 where load had no effect when the target and inducer orientations were easily distinguishable) and in Experiment 5 the effect of the mask was reduced (compared to Experiment 7 where load again had no effect when target and inducer were clearly separate). Thus, it appears that high perceptual load removes the impact of conscious (but not unconscious) distractors, and thereby reduces the perceived tilt-illusion under conscious compared to unconscious conditions. This suggestion is in accord with studies that found reduced suppression of weak, irrelevant distractors that were not consciously perceived compared to presentation of clearly visible distractors (e.g.

Tsushima, Sasaki & Watanabe, 2006; Tsushima, Seitz & Watanabe, 2008). According to this interpretation, the stimulus strength of distractors must be above a certain threshold in order to be actively suppressed by the attention system. The reduced ability of the visual system to distinguish signal from noise under high load (evident in Chapter 2), therefore, increased the impact of the noise (the context) under high load, specifically when the noise was presented subthreshold and therefore could not be actively suppressed.

Chapter 4

The role of perceptual load in pattern masking: Orientation detection and discrimination

4.1 Chapter Introduction

Experiments in the previous chapters have focussed on the spatial integration of orientation signals under perceptual load. However, as perception is a dynamic process that evolves over time (VanRullen & Thorpe, 2001), expanding this investigation into the temporal domain is a natural extension of the previous work. If high perceptual load reduces the efficiency to separate signal from noise, this is likely to also occur when signal and noise are presented successively. The current empirical chapter examines how the perceptual consequences of load evolve over time. More specifically, the experiments made use of the same physical stimuli as those used for the targets and noise masks in Chapter 2 where the noise mask was presented together with the target. However, instead of presenting target and mask together, here they were separated in time by varying the length of stimulus onset asynchronies (SOAs) between target and mask.

The literature on pattern backward masking (where target and mask overlap spatially) distinguishes between two types of mechanisms that can both contribute to the perceptual degradation of the target (Breitmeyer, 1984; Enns & DiLollo, 1997; Ganz, 1975; Kahneman, 1968; Spencer & Shuntich, 1970; Turvey, 1973). The following experiments investigate which of these mechanisms is likely to underlie the effects of perceptual load on target degradation in detection and discrimination and aim to establish the locus of these effects to early or later stages of visual processing.

According to the *integration masking* account, pattern masking strictly depends on the physical attributes of target and mask (e.g. mask contrast) and is thought to be unaffected by increases in display set size (increasing the number of possible targets) or other attentional manipulations. Here, degradation of the stimulus signal (the target) is due to the addition of external noise (the mask), much like the summation of target and mask features when they are presented together. Due to the poor temporal resolution of the visual system, this type of masking extends to short SOAs (of up to around 100ms) before performance starts to recover, but is always most pronounced when mask and target are presented together (SOA = 0; Figure 4.1a). Masking effects that extend to longer SOAs (e.g. up to 300ms SOA) have been attributed to attentional interactions, and more precisely to the *interruption* of central attentional mechanisms that were thought to be involved in transferring information from iconic visual memory to non-visual memory stores (e.g. working memory; Neisser, 1967; Sperling 1960; 1963; 1967). The idea was that once information had been stored in a non-visual form, it would be immune to visual masking (Michaels & Turvey, 1979). Therefore, the effects of *interruption* masking are strongest when the mask follows the target compared to when they are presented together (Figure 4.1b).

Figure 4.1

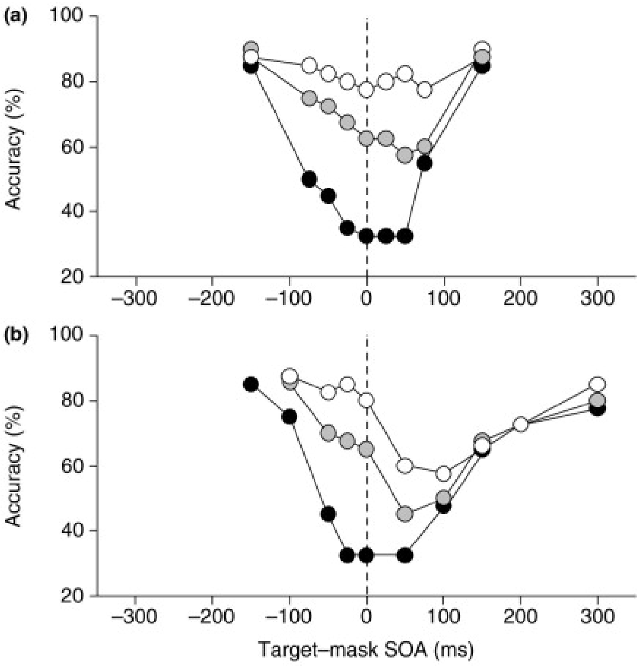

Figure 4.1. The effect of mask contrast and set size in pattern masking. (a) A single target letter was presented at 1 of 12 possible locations in a circular array. **(b)** 12 letters were presented and the location of the mask indicated the target letter. Varying mask contrast (white circles = low; grey circles = medium; black circles = high contrast) had its strongest effect at the 0ms SOA, while varying set size (1 vs. 12) most strongly modulated performance at SOAs longer than 0ms. (Spencer & Shuntich, 1970).

The mechanism of *interruption* masking is similar in spirit to the more recent proposal of *object substitution* masking (Di Lollo et al., 2000; Enns & Di Lollo, 1997) based on cortical re-entry activity (e.g. Lamme, 2000; Zeki, 1993). Here recurrent feedback among different levels of visual processing is assumed to be necessary to obtain a stable percept before information can be transferred to verbal memory (i.e. accessed by consciousness; Block 1995; Lamme, 2000; Lamme et al., 2002). Masking occurs due to the mismatch of information from high levels to lower levels relating to the target while at lower levels processing has already proceeded to the mask (Lamme et al., 2002).

The following experiments examined whether perceptual load increases the effectiveness of a pattern mask and if so which of the two proposed mechanisms is likely to underlie the observed effect. Furthermore, differences in the pattern of results for load effects on orientation detection versus discrimination may reveal the locus for such effects in the visual system.

4.2 Experiment 8

The purpose of Experiment 8 was to investigate how target detection in pattern masking is modulated by perceptual load in an unrelated task. Importantly, the design of the current experiment contrasts with previous attentional manipulations in pattern masking by minimizing spatial and decision uncertainty of the target location and additional sources of external noise from irrelevant distractors present in the usual set size manipulations (e.g. Averbach & Coriell, 1961; Enns & Di Lollo, 1997; Spencer, 1969; Spencer & Shuntich, 1970; Tata, 2002). Others have shown in detail how the behavioural benefits attributed to attention, when comparing single and multi-element displays, can be fully explained in terms of spatial uncertainty (Baldassi & Burr, 2000; Cameron, Tai, Eckstein, & Carrasco, 2004; Kinchla, 1992; Morgan et al., 1998; Shiu & Pashler, 1994; Solomon et al., 1997). According to these models, increasing the number of distractors (or empty target locations), increases the amount of noise that can interfere with identification of the target signal. Furthermore, the additional samples from distractor locations are thought to weaken the perceptual representation of the target, making it more vulnerable to interference by a subsequent mask.

In the following experiments the target was always presented at the same location (the centre of the screen) and had the same duration and onset timing on each trial throughout the experiment. Participants fixated their gaze at the location of the target throughout each trial, while perceptual load was modulated in a letter search task presented simultaneously in a circular array around fixation (the same letter search task used in the experiments of the previous chapters).

4.2.1 Method

Observers. Twenty observers (14 females, 18 to 25 years old, *Median* = 20) participated in Experiment 8. All were recruited through the University College London subject pool and had normal or corrected to normal vision and were naïve as to the purpose of the experiment. None of the observers had participated in any of the previous experiments.

Apparatus & Stimuli. All stimuli were created using MatLab (v2010a, The MathWorks, Nattick, MA) and Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) and presented on a 21.5” Monitor (1920 x 1080 pixel resolution, 60 Hz refresh rate) in a dimly lit room. A chin and forehead rest was used to fix viewing distance to 57cm. All stimuli were presented on a medium grey background. Participants performed a dual-task, detecting a masked target at fixation while a letter search task was used to vary the level of perceptual load. For the letter search task, eight, black letters (subtending 0.6° x 0.9° of visual angle) were presented equidistant from fixation each centred at 3.2° eccentricity. In the low load condition participants searched for a Z or N among seven Os, whereas in the high load condition participants searched for a Z or

N among seven different, angular letters (randomly selected from K, E, F, H, L, M, T, W, Y or Z on each trial). On a random half of the trials the presented target letter was Z and otherwise N. The target for the detection task was a vertical Gabor patch (1.5° in diameter, 32% Michelson contrast, 6 cpd) presented at fixation, while the mask consisted of an orientation band-passed noise patch with mean orientations symmetrically centred on 45° clockwise and 45° counter-clockwise from vertical (4.8° in diameter, 45% rms contrast, 10° orientation bandwidth, low-pass spatial frequency filtered with a 10 cpd cut-off). The mask was presented at fixation, spatially overlapping the target location. Figure 4.2 illustrates the stimulus display under high load.

Figure 4.2

Figure 4.2. Schematic illustration of example displays in the high load condition of Experiment 8. Participants searched for a Z or N among seven angular letters (high load) or seven Os (low load, not shown) while also detecting a vertically oriented target at the centre. **(left panel)** Target in isolation. **(middle panel)** Mask in isolation. **(right panel)** Target and mask together (i.e. SOA = 0ms).

Procedure. At the beginning of each trial a fixation point was presented at the centre of the screen for 700ms. Following a blank interval (500ms) the central target was presented together with the letter search task. While the target disappeared after 17ms, the letter search task remained on screen for 167ms in total. After a variable delay of 0, 50, 67, 83 or 150ms (or no mask) after target onset (SOA) the mask was presented for 33ms. Observers then responded first to the letter-search task (indicating the target letter "Z" or "N" by pressing one of two designated keys on the left side of the keyboard) and then to the central target ("present" or "absent", pressing one of two keys on the right side of the keyboard).

4.2.2 Results

Visual search. Under high perceptual load, visual search accuracy (*M* = 81%, *SEM* = 2.13) was significantly reduced and reaction time (*M* = 1153ms, *SEM* = 21) significantly increased compared to accuracy (*M* = 96%, *SEM* = .64) and reaction time (*M* = 876ms, *SEM* = 28) under low load; $t(19) = 8.06$, $p < .001$ and $t(19) = 10.66$, $p < .001$, for the load effects on accuracy and reaction time, respectively. This confirms the effectiveness of the attentional manipulation. Performance in the letter search task was comparable to that in previous experiments.

Orientation target detection

Accuracy. To test whether the mask was effective in reducing target detection, accuracy scores were subjected to a two-factor, within subject, repeated measures ANOVA with the factors Mask (mask vs. no mask) and Load (low vs. high). This showed a main effect of Mask, with detection accuracy significantly worse in the presence of the pattern mask (*M* = 68%, *SEM* = 2.2) compared to the no mask condition (*M* = 82%, *SEM* = 2.1); $F(1,19) = 23.99$, $p < .001$, $\eta_p^2 = .56$. Furthermore, high perceptual load reduced target detection overall (*M* = 66%, *SEM* = 2.4) compared to low load (*M* = 84%, *SEM* = 1.7), *F*(1,19) = 48.59, $p < .001$, $\eta_p^2 = .72$. The main effect of Mask indicates that masking was effective in reducing target detection, irrespective of the level of perceptual load in the visual search task. Conversely, the main effect of Load indicates that high perceptual load in the visual search reduced target detection at fixation regardless of the presence or absence of the pattern mask. When comparing the effects of load and mask in this way, no significant interaction was found as the overall effect of presenting a mask was comparable to the overall effect if increasing the level of load. Therefore, the effect of load must be assessed across the different SOAs between target and mask.

Figure 4.3

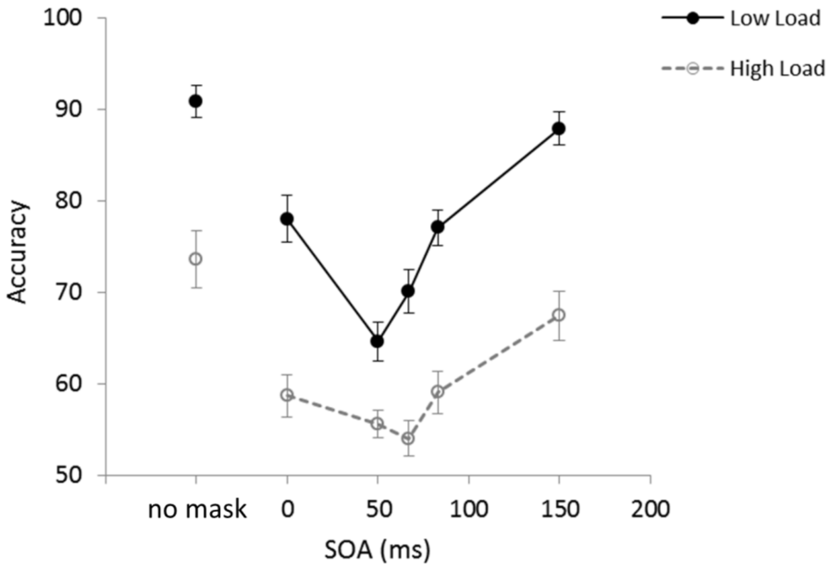

Figure 4.3. Detection accuracy for the target under low and high perceptual load in the letter search task in Experiment 8. Accuracy is shown as a function of SOA (including a no mask control condition). Masking was stronger under high perceptual load and started to

recover slightly later (around 76ms SOA) compared to the low load condition (around 50ms). The level of perceptual load also had a strong effect on target detection when no mask was presented (no mask condition). Data were averaged over 20 participants and error bars represent ±1 *SEM*.

Figure 4.3 illustrates the reduced detection accuracy under high compared to low perceptual load across the different SOAs, as well as, when no mask was present. To analyse how perceptual load modulated target detection across the SOAs a two-factor, within subject repeated measures ANOVA with the factors Load and SOA was performed. This revealed significant main effects of Load, $F(1,19) = 61.22$, $p < .001$, $\eta_p^2 = .76$ and SOA, $F(2.72,51.72) = 39.54$, Greenhouse-Geisser corrected, $p < .001$, $\eta_p^2 = .68$. In addition, the dependency of masking on the level of perceptual load was indicated by a significant interaction between Load and SOA, $F(3.45,64.91) = 5.00$, Greenhouse-Geisser corrected, $p < .01$, $\eta_p^2 = .21$.

The results provide evidence for the hypothesis that high perceptual load in an unrelated task increases the effectiveness of a pattern mask. Moreover, the interaction between Load and SOA may indicate a slower recovery from masking under high compared to low load. This was confirmed by pairwise comparisons (with Bonferroni adjustment for multiple comparisons) showing that under low load, detection accuracy significantly improved from 50ms SOA ($M = 65\%$, $SEM = 2.2$) to 67ms SOA ($M = 70\%$, $SEM = 2.5$), $t(19) = 3.83$, $p < .05$; whereas under high load performance did not significantly differ between these SOAs (50ms $M = 56\%$, $SEM = 1.5$; 67ms $M = 54\%$, $SEM = 2.0$), $t(19) = .82$, $p = 1.00$. Lastly, only under low load did performance at the longest SOA (150ms $M = 88\%$, $SEM = 1.9$) fully recover to the level of the no mask condition ($M = 91\%$, $SEM = 1.8$), $t(19) = 2.21$, $p = .59$; while under high load performance at this SOA ($M = 66\%$, $SEM = 2.8$) was still significantly suppressed, $t(19) = 3.48$, $p < .05$, compared to the no mask condition ($M = 74\%$, $SEM = 3.3$).

Orientation target detection

Sensitivity. Since observers in this experiment reported the presence or absence of the target, simple analysis of accuracy scores may not reflect true sensitivity to the target stimulus. For example, it is possible that under high perceptual load, observers were more likely to confuse the target with the mask, thus responding “present” more often than “absent” compared to the low load condition. In order to control for any differences in response criterion between the two conditions, hit and false alarm rates were calculated for each SOA, load level and observer to calculate d’ scores (Green & Swets, 1966).

The same, within-subject, repeated measures ANOVA with the factors Load and SOA was conducted on the d’ scores, which showed again significant main effects of Load, $F(1,19) = 98.29$, $p < .001$, $\eta_p^2 = .81$, and SOA, $F(3.12,59.32) = 43.20$, Greenhouse-Geisser corrected, $p < .001$, $\eta_p^2 = .70$, as well as, a significant interaction, $F(5,95) = 4.60$, $p < .01$, $\eta_p^2 = .20$ (Figure 4.4).

The d’ scores revealed comparable masking curves under low and high load to those obtained from the accuracy scores. Masking peaked at SOAs greater than 0ms and was more pronounced under high compared to low perceptual load. Furthermore, masking recovered more slowly under high compared to low perceptual load.

Figure 4.4

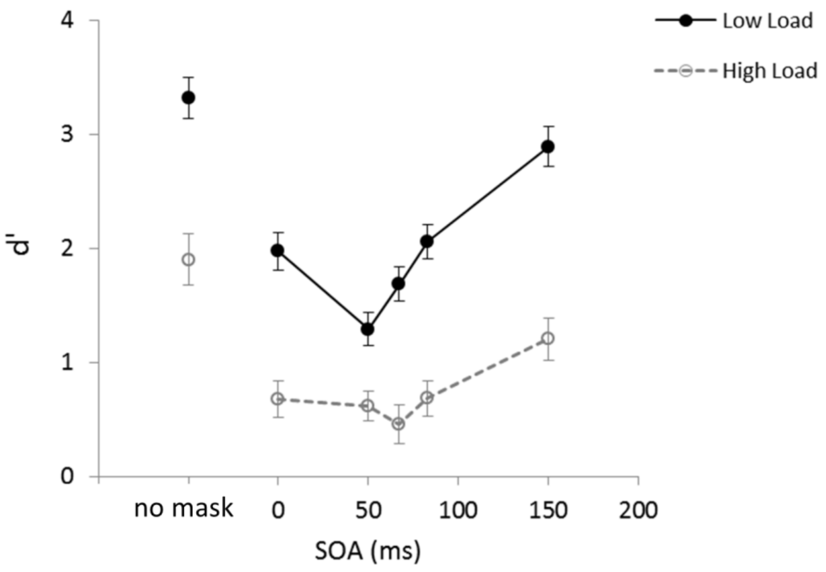

Figure 4.4. d' (d-prime) scores, indicating detection sensitivity to the oriented target under low and high perceptual load in the letter search task in Experiment 8. Data were averaged over 20 participants and error bars represent ±1 *SEM*.

Pairwise comparisons (with Bonferroni adjustment for multiple comparisons) showed that while under low load target detection sensitivity (d') significantly improved from 50ms SOA (*M* = 1.3, *SEM* = .15) to 83ms SOA (*M* = 2.1, *SEM* = .16), *t*(19) = 4.83, *p* < .01, under high load there was no change in sensitivity from 50ms SOA (*M* = .62, *SEM* = .14) to 83ms SOA (*M* = .68, *SEM* = .16); *t*(19) = .47, *p* = 1.00. In contrast the improvement from 83ms SOA to 150ms SOA was significant in both load conditions; *t*(19) = 6.89, *p* < .001 and *t*(19) = 3.69, *p* < .05, for the increases in sensitivity under low and high load, respectively.

4.2.3 Discussion

The results of Experiment 8 demonstrate that the effectiveness of pattern masking is indeed increased under high perceptual load. Moreover, the masking curves indicate that masking was strongest at SOAs greater than 0ms (target and mask presented together) for both load conditions. However, this effect increased under high perceptual load, where masking strength peaked at 67ms SOA, compared to 50ms SOA under low load.

Together, the results suggest that a form of *interruption masking* is more likely to mediate the effects of load compared to an account in terms of *integration masking*. First of all, masking was most pronounced at SOAs greater than zero for both load conditions. An account of pattern masking in terms of pure *integration masking* (which assumes that masking occurs due to the increase of external noise in the target representation and the poor resolution of the visual system at short SOAs), however, would predict masking strength to be maximal at 0ms SOA. Due to the specific properties of the mask in this experiment (containing orientations relatively dissimilar from the target) the target was predicted to be visible on some of the trials where target and mask were presented together (SOA = 0). However, this does not predict a further decrease in performance at SOAs greater than zero (backward masking).

Secondly, the masking curves obtained in the two load conditions differed primarily after the 50ms SOA, with performance recovering under low load while deteriorating further under high load. This suggests that load modulated masking strength independent of local integration occurring at 0ms SOA.

4.3 Experiment 9

The purpose of Experiment 9 was to assess how perceptual load modulates pattern masking for orientation discrimination (in contrast to orientation detection in Experiment 8). It is plausible to assume that detection of an oriented target among oriented noise is less computationally demanding as discriminating the orientation of such a target. Indeed, it has been shown repeatedly that identifying a target within an array of distractors requires serial processing when target and distractors share common features (discrimination), whereas target detection appears to occur automatically when no features are shared among target and distractors (detection) (e.g. Bergen & Julesz, 1983; Treisman & Gelade, 1980). Moroever, this also occurs when different tasks (e.g. detection vs. discrimination) are performed on targets that only vary in a single feature such as line orientation (Sagi & Julesz, 1985). If perceptual load affects additional processes involved with discrimination, the effects of load should be evident at later SOAs compared to detection. However, if the pattern obtained with orientation discrimination does not differ from orientation detection, perceptual load must influence the earlier of the two components as modulation at the early stage of low level signal accumulation would result in equally reduced performance for detection and discrimination.

4.3.1 Method

Observers. Twelve new observers participated in Experiment 9 (5 females, 20 to 31 years of age, *Median* = 24). All were recruited through the University College London subject pool, had normal or corrected to normal vision and were naïve to the purpose of the experiment.

Stimuli & Procedure. All aspects of stimulus timing (SOAs, target and mask durations) and stimulus dimensions (size, contrast, spatial frequency) were identical to those in the previous experiment. However, instead of presenting the target on 50% of the trials, it was presented on every trial and was either oriented 20° clockwise or 20° counter-clockwise from vertical. Instead of detecting the presence of the target, observers were asked to discriminate its orientation, using the left (for counter-clockwise from vertical) or right (for clockwise from vertical) arrow keys on the right side of the keyboard.

The same letter search task was used as in the previous experiment to modulate the level of perceptual load. Observers responded first to the letter search task (pressing one of two designated keys on the left side of the keyboard to indicate the target letter) before making the orientation judgment.

4.3.2 Results

Visual search. Under high perceptual load, letter search accuracy (*M* = 74%, *SEM* = 3.5) was significantly reduced, while reaction time (*M* = 1197ms, *SEM* = 60) was significantly increased compared to accuracy (*M* = 93%, *SEM* = 2.1) and reaction time (*M* = 966ms, *SEM* = 65) under low load; $t(11) = 7.62$, $p < .001$ and $t(11) = 8.96$, $p < .001$ for accuracy scores and reaction times respectively.

Orientation target discrimination. A two-factor repeated measures ANOVA with the factors Mask (mask vs. no mask) and Load (high vs. low) showed a main effect of Mask, $F(1,11) = 17.60$, $p < .01$, $\eta_p^2 = .62$; indicating that masking was effective in reducing target orientation discrimination accuracy, irrespective of the level of load in the visual search task. Average discrimination accuracy in the presence of the mask was 77% correct (*SEM* = 4.0) and 89% correct (*SEM* = 2.8) when no mask was presented. Furthermore a significant main effect of Load, $F(1,11) = 34.40$, $p < .001$, $\eta_p^2 = .76$, suggested that overall discriminability was reduced under high ($M = 76\%$, $SEM = 3.7$) compared to low load ($M = 90\%$, $SEM = 3.0$).

To assess modulation by perceptual load of the effectiveness of the mask across the different SOAs, a two-factor repeated measures ANOVA was performed with the factors Load and SOA. This revealed significant main effects of Load, $F(1,11) = 50.28$, $p < .001$, $\eta_p^2 = .82$, and SOA, $F(5,55) = 24.45$, $p < .001$, $\eta_p^2 = .69$, as well as, a significant interaction between Load and SOA; $F(5,55) = 2.49$, $p < .05$, $\eta_p^2 = .19$.

Figure 4.5

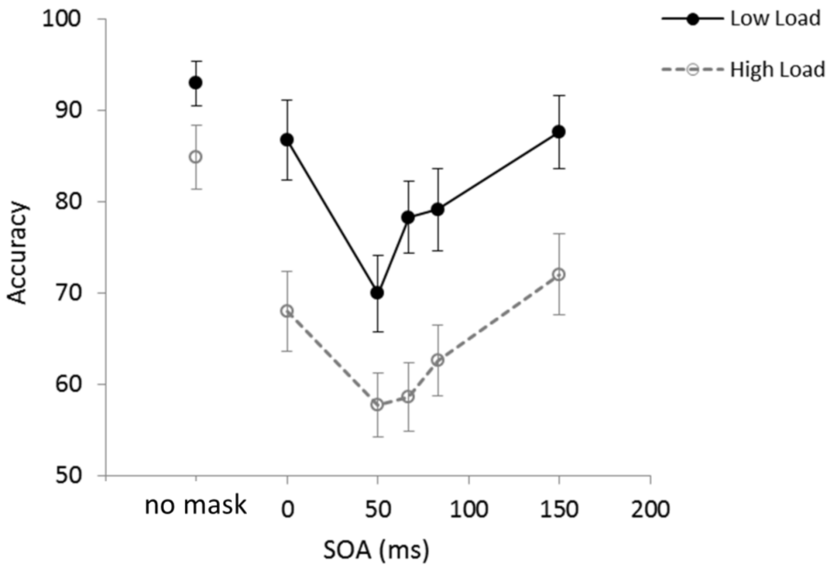

Figure 4.5. Orientation discrimination accuracy for the target under low and high perceptual load in the letter search task in Experiment 9. Accuracy is shown as a function of SOA (including the no mask condition). Masking was stronger under high perceptual load

and started to recover slightly later (around 67ms SOA) compared to the low load condition (around 50ms). Data were averaged over 12 observers and error bars represent ±1 *SEM*.

The main effect of Load indicates that high perceptual load significantly reduced accuracy in the orientation judgment across the SOAs, making the pattern mask more effective compared to low load. Moreover, the interaction between Load and SOA demonstrates that mask effectiveness was differently modulated across the SOAs under high compared to low load, suggesting a slower recovery from mask suppression under high load.

Indeed, pairwise comparisons (with Bonferroni adjustment for multiple comparisons) showed that under low load performance significantly improved from 50ms SOA (*M* = 70%, *SEM* = 4.2) to 67ms SOA (*M* = 78%, *SEM* = 3.9), *t*(11) = 3.78, *p* < .05, while under high load there was no significant difference in performance between the 50ms (*M* = 58%, *SEM* = 3.5) and 67ms (*M* = 59%, *SEM* = 3.8) SOAs; *t*(11) = .58, *p* = 1.00. Figure 4.5 illustrates how performance starts to recover after 50ms SOA under low load, while it remains suppressed at least until 67ms SOA under high load. Lastly, while at the longest SOA of 150ms, performance under low load (*M* = 88%, *SEM* = 4.0) had fully recovered to the level of the no mask condition (*M* = 93%, *SEM* = 2.5), *t*(11) = 2.41, *p* = .52, this was not the case under high load (150ms SOA *M* = 72%, *SEM* = 4.4, no mask *M* = 85%, *SEM* = 3.5), *t*(11) = 3.95, *p* < .05.

4.3.3 Discussion

The pattern of results from Experiment 9 is very similar to the previous experiment on detection. The presentation of the pattern mask at the same SOAs produced a U-shaped masking function under both load conditions with maximal masking occurring at around 50ms SOA. High perceptual load in the visual search task caused an overall increase in the effectiveness of the pattern mask to reduce orientation discrimination accuracy. However, as in the previous experiment, masking extended to longer SOAs under high compared to low load. This suggests that perceptual load does not simply increase the magnitude of masking independent of target mask SOA, but prolongs the time period in which the target is vulnerable to interference from the mask. Moreover, the similar pattern obtained in the two experiments suggests that perceptual load modulates an early stage in low level visual integration and does not further degrade target information at later stages involved in discrimination.

4.4 Chapter Conclusions

The results reported in this chapter provide evidence for an effect of perceptual load on the masking functions for orientation detection and discrimination using pattern masks. Using a mask which was noise band-passed in the orientation domain (with mean noise orientations centred at 45° clockwise and 45° counter-clockwise from vertical) its maximal masking effect occurred around 50ms SOA under low and high load. This finding is in line with previous studies using oriented masks (e.g. Wehrhahn, Li & Westheimer, 1996) without manipulating attentional factors. Moreover, it suggests that perceptual load effects rely on mechanisms of *interruption masking* and are not simply due to sensory integration as if target and mask are presented together (SOA = 0ms) suggested by *integration masking* theories. Although usually masking by noise or structure produces masking functions where performance recovers monotonically as SOA increases from 0ms onwards (Kinsbourne & Warrington, 1962; Scharf & Lefton, 1970; Shiller, 1966), U-shaped backward masking functions (when masking strength varies non-monotonically at SOAs greater than 0ms) have been observed specifically when target and mask have the same intensities (e.g. same contrast and duration; Michaels & Turvey, 1973; Purcell & Stewart, 1970; Spencer & Shuntich, 1970, Turvey, 1973; Weisstein, 1971), as was the case in the current experiments. However, an increase in mask intensity has not been shown to extend masking to longer SOAs (as was the case for high perceptual load in both experiments). This finding implies that for both, orientation detection and discrimination, perceptual load modulates target and mask integration via a process that appears independent of low level mask properties (at least those relating to mask and target intensity).

Most of the models that aim to explain backward masking in the visual system can be divided into two categories. The first category comprises models of pure feedforward processes where the mask inhibits feedforward activity from the target. This occurs before target processing is completed and target information is made available to other, non-visual forms of memory. Inhibition occurs because the transient onset of the mask inhibits the activity in slower channels involved in target identification (e.g. Breitmeyer, 1984). The other category of models consists of those that argue pattern masking depends on interference in recurrent feedback processing. These theories are based on findings that suggest recurrent activity (the flow of information from lower to higher levels and back) is a necessary condition for conscious perception (e.g. Enns & DiLollo, 2000; Lamme 2000; Lamme & Roelfsema, 2000; Lamme et al., 2002). Masking occurs in these models due to the disruption of recurrent processes because when information from higher levels of processing is fed back to lower levels, the information content of the target at the lower levels has been replaced by the mask. Evidence for recurrent feedback models of pattern masking has been obtained in a study that investigated the effect of pattern masking on figure ground segregation (Lamme et al., 2002). The authors based their experiments on the finding that while the orientation of a line segment is signalled in V1 at latencies as early as 55ms and pure feedforward mechanisms are sufficient to produce orientation selectivity (e.g. Ferster, Chung & Wheat, 1996; Miller & Ferster, 2000), signalling of figure-ground segregation in V1 occurs at later latencies (around 100ms) and is likely to involve lateral interactions and feedback from extrastriate areas (Lamme, Super, & Spekreijse, 1998). Using single cell recordings in monkeys, the authors found that pattern masking specifically reduced responses at those SOAs where activity in V1 signals figure ground segregation while earlier information relating the orientation of the display elements was not affected.

This finding is in line with the results reported here (as well as those from previous chapters showing load effects on contextual modulation) where the strongest effect of the level of load in modulating the masking curve was observed at SOAs beyond 50ms. The question of whether load also affects masking effects that are thought to occur at higher levels of object recognition (e.g. when mask and target do not spatially overlap) is explored in the next chapter.

Chapter 5

The role of perceptual load in metacontrast masking and object-substitution

5.1 Chapter Introduction

When a briefly presented target is followed by a mask whose contours are directly adjacent to that of the target without spatially overlapping, target visibility can be drastically reduced. This phenomenon represents another type of backward masking, called *metacontrast masking*. Traditionally, empirical attempts to uncover the underlying mechanisms of metacontrast masking have focussed on the physical characteristics of target and mask stimuli, such as intensity and contour proximity. These types of investigations resulted in theoretical models of metacontrast masking that sought to explain the phenomenon in terms of local inhibitory interactions among neurons representing adjacent contours in low level vision (summarized in Chapter 1).

However, a number of studies challenged the traditional view of metacontrast masking by demonstrating that higher level cognitive processes such as perceptual grouping modulate the masking function (Caputo, 1998; Ramachandran & Cobb, 1995) and others have claimed a role for visual selective attention (Boyer & Ro, 2007; Enns & Di Lollo, 1997; Ramachandran & Cobb, 1995; Tata, 2002; Tremblay & Mack 1999). However, most of the studies investigating attentional effects on metacontrast masking either did not directly manipulate attention (e.g. correlating resistance to *inattentional blindness* of some meaningful stimuli to resistance in metacontrast masking, Tremblay & Mack 1999; or presuming attentional involvement in perceptual grouping, Ramachandran & Cobb, 1995), simply increased the spatial uncertainty of the target by making it part of a larger display (e.g. Averbach & Coriell, 1961; Enns & Di Lollo, 1997; Tata, 2002), or altered the duration of the mask with respect to the duration of the target (Di Lollo et al., 2000). Moreover, the role of perceptual load and restriction of available attentional resources in metacontrast masking has not been explored.

The current chapter first investigates the effect of perceptual load on metacontrast masking (Experiments 10-12) using a similar paradigm as described in the previous chapter on pattern masking. Here, the target location and target onset timing is fixed, eliminating effects of spatial uncertainty, while perceptual load is varied in a completely unrelated letter search task. The results point to a role of perceptual load in metacontrast masking, specifically at later SOAs where classical pattern masking is no longer influenced by load (Chapter 4). Strong reductions in target discriminability at late SOAs (i.e. beyond 167ms or even 300ms SOA) under high perceptual load suggest an underlying mechanism in line with *object substitution* theories. Therefore, the final part of the chapter directly investigates the effects of load on *object substitution* masking (Enns & Di Lollo, 1997; Di Lollo et al., 2000). Here, the use of a mask consisting of only four small dots surrounding the target location distinguishes local, low level processes thought to be involved in metacontrast masking from higher level processes, presumably relying on recurrent feedback to maintain a stable percept of the target.

5.2 Experiment 10

The purpose of Experiment 10 was to establish the effects of perceptual load on metacontrast masking. More specifically, based on previous research claiming attentional effects on metacontrast masking (Ro & Boyer, 2007; Tata, 2002) and the results of the previous chapter indicating stronger pattern masking under load, I hypothesized that high perceptual load in an unrelated task should increase the effectiveness of a metacontrast mask.

Furthermore, the effects of high perceptual load should be observable at later SOAs between target and mask, compared to low load. That is, if *metacontrast* masking under load is indeed the result of *object substitution* masking due to re-entrant feedback activity, limiting perceptual resources under high load should increase the time needed to process the target, making it more vulnerable to interference from the mask at later SOAs. Therefore, metacontrast masking should not only be more effective under high perceptual load, but should also lead to a slower recovery of target discrimination accuracy, shifting the period of maximal masking to longer SOAs. Importantly, perceptual load was manipulated in a letter search task that was always presented in the same clock face configuration around fixation. Thus, the search task was spatially separated from the shape discrimination task at the centre. Furthermore, the central target shape and mask locations, their durations, as well as target onset timing were identical on each trial. Lastly, unlike in previous studies, target set size was not varied (only a single masked target was presented on each trial) and the attentional manipulation was solely a result of target-distractor similarity in the letter search task.

This design allowed for directly assessing the effects limiting attentional resources for target processing in metacontrast masking without relying on factors that may or may not have engaged attentional mechanisms in previous experiments (such as perceptual grouping, spatial uncertainty of the target location or variability in mask and target durations; discussed in more detail in Chapter 1).

5.2.1 Method

Observers. Fourteen observers (11 females, 19 to 30 years old, *Median* = 24) participated in the experiment. All were recruited through the University College London subject pool, had normal or corrected to normal vision and were naïve to the purpose of the experiment. None of the observers had participated in any of the previous experiments.

Apparatus & Stimuli. All stimuli were created using MatLab (v2007a, The MathWorks, Nattick, MA) and Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) and presented on a 21.5” Monitor (1920 x 1080 pixel resolution, 60 Hz refresh rate). A chin and forehead rest was used to fix viewing distance to 57cm. Stimuli were presented in light grey on a black background. Participants performed a dual task, searching a circular array for a target letter while discriminating a shape at fixation. The letter search task consisted of eight, light grey letters (subtending 0.6° x 0.9° of visual angle) presented equidistant from fixation, each centred at 3.2° eccentricity. In the low load condition participants searched for a Z or N among seven Os, whereas in the high load condition participants searched for a Z or N among seven different, angular letters (randomly selected from K, E, F, H, L, M, T, W, Y or Z on each trial). On a random half of the trials the presented target letter was Z and otherwise N.

The central shape discrimination task consisted of a single diamond or square shape (1° of visual angle) briefly presented at fixation. On a random half of the trials within a block the shape was a diamond, on the other half it was a square. The shape was masked after a variable delay (or no mask was presented). The mask (2.8° in diameter) had the shape of a round disc with a central cut-out that matched both, the contours of the diamond and the square. It was fitted tightly around the contours of the two targets without spatially overlapping them (Figure 5.1).

Figure 5.1

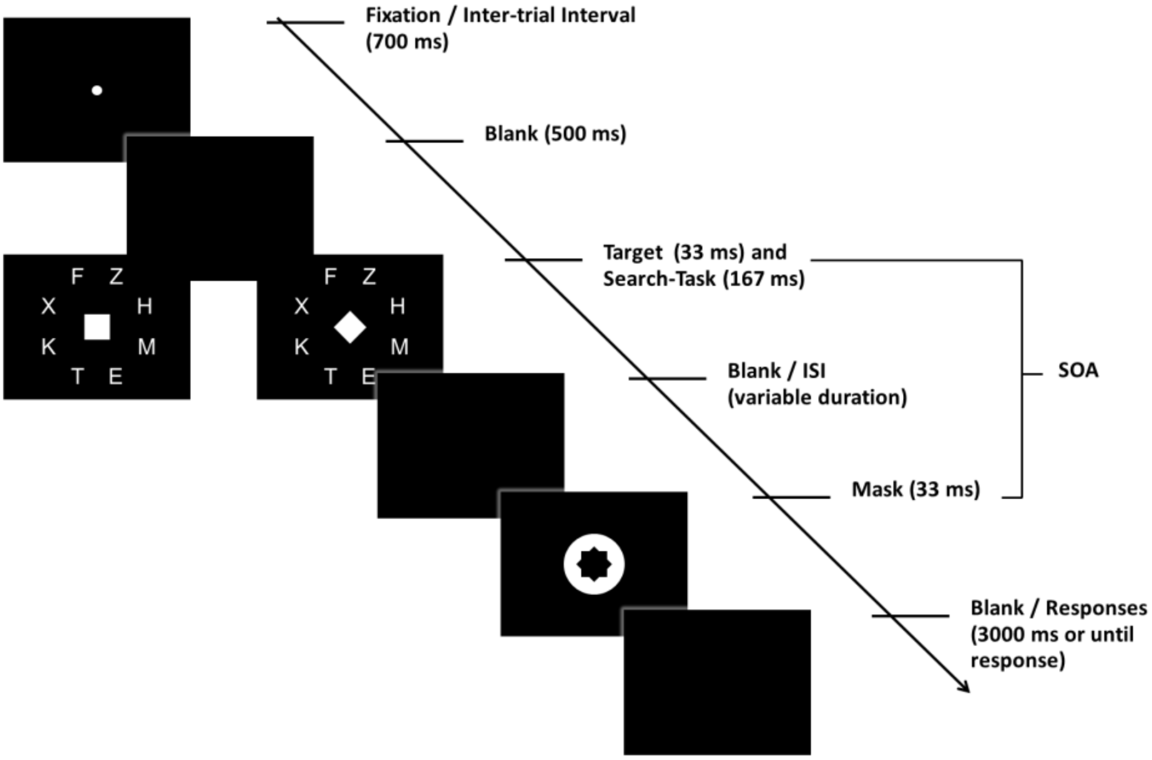

Figure 5.1. Visual discrimination task with metacontrast masking and letter search task in the high load condition of Experiment 10. The stimuli were presented in light grey on a black background. The mask tightly fits around part of the contour of the target without overlapping with the target spatially. Target and letter search task were always presented together. In the low load condition all non-target letters were Os (not shown). After the presentation of the mask participants were first asked to decide whether a Z or N had been presented. Then they had to indicate whether they had seen a diamond or a square. Shown in the brackets are the durations of each stimulus.

Procedure. At the beginning of each trial a fixation point was presented at the centre of the screen for 700ms. Following a blank interval (500ms) the central target (diamond or square) was presented together with the letter search task. While the central target disappeared after 33ms, the letter search task remained on screen for 167ms in total. After a variable delay of 50, 167, 317 or 500ms after target onset (SOA) the mask (or a blank screen in the no mask condition) was presented, also for 33ms. Observers then responded first to the letter search task (indicating the target letter "Z" or "N" by pressing one of two designated keys on the left side of the keyboard) and then to the central target ("diamond" or "square", pressing one of two keys on the right side of the keyboard).

5.2.2 Results

Visual search. As expected, in conditions of high load, visual search accuracy (*M* = 69%, *SEM* = 2.80) was significantly reduced and reaction time (*M* = 1025ms, *SEM* = 77) significantly increased compared to the accuracy (*M* = 91%, *SEM* = 1.35) and reaction time (*M* = 852ms, *SEM* = 51) under low load; $t(13) = 10.59$, $p < .001$ and $t(13) = 4.36$, $p < .001$ for the load effects on accuracy and reaction time, respectively; confirming the efficacy of the attentional manipulation.

Shape discrimination. Overall accuracy in the shape discrimination task was significantly reduced in the presence of the mask (*M* = 59%, *SEM* = 2.3) compared to the no mask condition (*M* = 82%, *SEM* = 3.4, $F(1,13) = 76.76$, $p < .001$, $\eta_p^2 = .86$; within subjects repeated measures ANOVA), confirming that the metacontrast mask effectively reduced target discriminability. The same ANOVA also showed a main effect of Load, $F(1,13) = 11.37$, $p < .01$, $\eta_p^2 = .47$; indicating overall reduced performance under high (*M* = 67%, *SEM* = 2.8) compared to low perceptual load (*M* = 74%, *SEM* = 2.7), irrespective of the presence of the mask.

Increasing perceptual load in the letter search task had a marked effect on the discriminability of the central target shape (Figure 5.2). A two-factor, within subject

ANOVA with the factors Load and SOA revealed significant main effects of Load, $F(1,13) = 33.30, p < .001, \eta_p^2 = .72$, and SOA, $F(4,52) = 39.80, p < .001, \eta_p^2 = .75$, as well as, a significant interaction between Load and SOA, $F(2.46,32.01) = 3.54, p < .05$ after Greenhouse-Geisser adjustment, $\eta_p^2 =.21$.

Figure 5.2

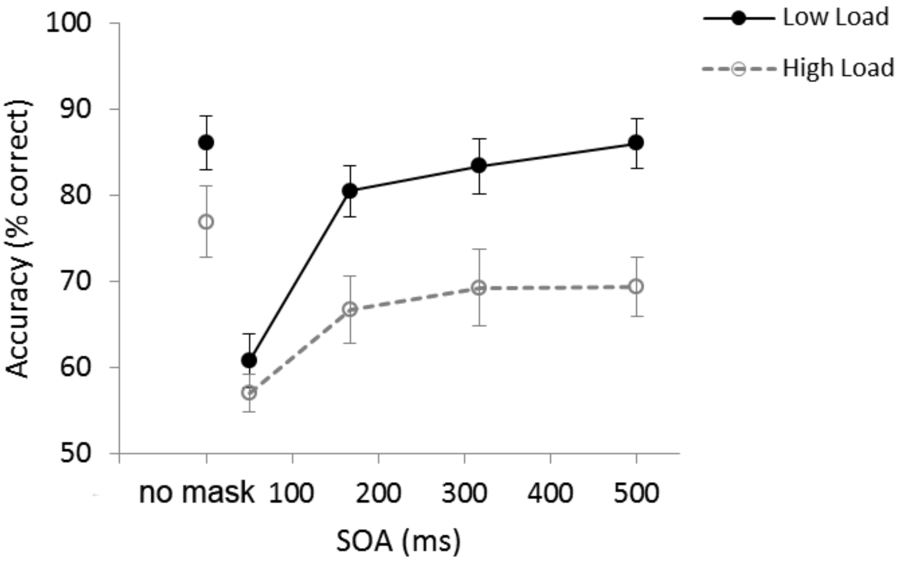

Figure 5.2. Discrimination accuracy for the masked target under low and high perceptual load in the letter search task in Experiment 10. Data were averaged over 14 participants and error bars represent ±1 *SEM*.

The main effect of Load indicates that high perceptual load significantly increased masking strength, whereas the main effect of SOA suggests that performance significantly improved with longer SOAs. The observed interaction between Load and SOA demonstrates that the strength of masking across the different SOAs depended on the type of load in the letter search task. Moreover, pairwise comparisons (with Bonferroni adjustment for multiple comparisons) showed that beyond 50ms SOA, performance under high perceptual load recovered more slowly than under low load. Accuracy in the low load condition was significantly lower at 50ms SOA (*M* = 61%, *SEM* = 3.1) compared to 167ms SOA (*M* = 81%, *SEM* = 3.0); *t*(13) = 5.43, *p* < .01. However, under high load there was no significant difference between the 50ms SOA (*M* = 57%, *SEM* = 2.2) and 167ms SOAs (*M* = 67%, *SEM* = 3.8); *t*(13) = 3.03, *p* = .10. On the other hand, comparing performance in the no mask condition to the 50ms SOA showed significant differences for both, low perceptual load, *t*(13) = 7.78, *p* < .001, and high perceptual load, *t*(13) = 5.99, *p* < .001. Finally, under low load, performance at 167ms SOA (*M* = 81%, *SEM* = 3.0) already recovered to the level of the no mask condition (*M* = 86%, *SEM* = 3.2); *t*(13) = 2.56, *p* = .24, while under high load performance at this SOA (*M* = 67%, *SEM* = 3.8) was still significantly suppressed by the mask (no mask *M* = 77%, *SEM* = 4.1); *t*(13) = 7.56, *p* < .001.

5.2.3 Discussion

The results suggest that high perceptual load in an unrelated task increases the effectiveness of metacontrast masking. Although under both levels of perceptual load, masking was maximal at 50ms SOA, the effect persisted under high load while performance under low load recovered to the level of the no mask condition immediately. This pattern suggests a mechanism of prolonged target processing under high load. While under low load, sufficient perceptual resources are available to process the target, the lack of such resources under high load makes the target more vulnerable to interference from the mask at longer SOAs.

5.3 Experiment 11

In Experiment 10, perceptual load in the letter search task appeared to influence masking strength even at the longer SOAs of 317ms and 500ms. Here, target discrimination accuracy was consistently lower under high compared to low load. However, it remains unclear if any significant recovery under low or high load occurred at SOAs beyond 167ms.

The purpose of Experiment 11 was to assess if under low load, performance truly recovered by 167ms SOA (reaching an asymptote), or if it continued to significantly improve at later SOAs. Furthermore, the experiment tested whether performance under high load showed any sign of improvement at later SOAs. To this end, Experiment 10 was repeated with the addition of an extra SOA at 400ms.

5.3.1 Method

Observers. Fourteen observers (8 females, 20 to 29 years of age, *Median* = 24), none of which had participated in the previous experiment, participated in Experiment 11. All were recruited from the University College London subject pool and had normal or corrected to normal vision.

Stimuli and Procedure. All aspects of the stimuli and procedure were identical to that of Experiment 10, except that an additional SOA was used at 400m. In addition, the 300ms SOA was changed to 317ms.

5.3.2 Results

Visual search. As previously, performance in the letter search task under high load (*M* = 78%, *SEM* = 2.68) was significantly reduced compared to low load (*M* = 94%, *SEM* = 1.64) and responses were significantly slower under high load (*M* =

1506ms, *SEM* = 99) compared to low load (*M* = 1200, *SEM* = 72); *t*(13) = 8.42, *p* < .001 and *t*(13) = 5.70, *p* < .001, for accuracy and reaction time data, respectively.

Shape discrimination. The mask was again effective in significantly reducing overall target discriminability (*M* = 59%, *SEM* = 3.0) compared to when no mask was present (*M* = 83%, *SEM* = 2.8); *F*(1,13) = 62.21, *p* < .001, η_p^2 = .83. Furthermore, high perceptual load significantly reduced target discrimination (*M* = 66%, *SEM* = 3.0) compared to low load (*M* = 76%, *SEM* = 2.7), irrespective of the presence of a mask, *F*(1,13) = 13.62, *p* < .01, η_p^2 = .51.

The data showed a similar pattern to Experiment 10, with masking being most effective at 50ms SOA under low and high load but recovering more slowly under high load (Figure 5.3). A two-factor, within subject repeated measures ANOVA with the factors Load and SOA revealed significant main effects of Load (*F*(1,13) = 52.39, *p* < .001, η_p^2 = .80) and SOA (*F*(5,65) = 23.48, *p* < .001 after Greenhouse-Geisser adjustment, η_p^2 = .64), as well as, a significant interaction between Load and SOA (*F*(5,65) = 3.34, *p* < .05 after Greenhouse-Geisser adjustment, η_p^2 = .20). Pairwise comparisons (with Bonferroni adjustment for multiple comparisons) for SOAs in each load condition showed that while accuracy significantly increased from 167ms (*M* = 80%, *SEM* = 2.6) to 500ms SOA (*M* = 91%, *SEM* = 2.4) under low load, *t*(13) = 5.72, *p* < .01, this was not the case for accuracy at 167ms (*M* = 68%, *SEM* = 3.2) and 500ms SOAs (*M* = 72%, *SEM* = 3.3) under high perceptual load, *t*(13) = 1.16, *p* = 1.00.

Figure 5.3

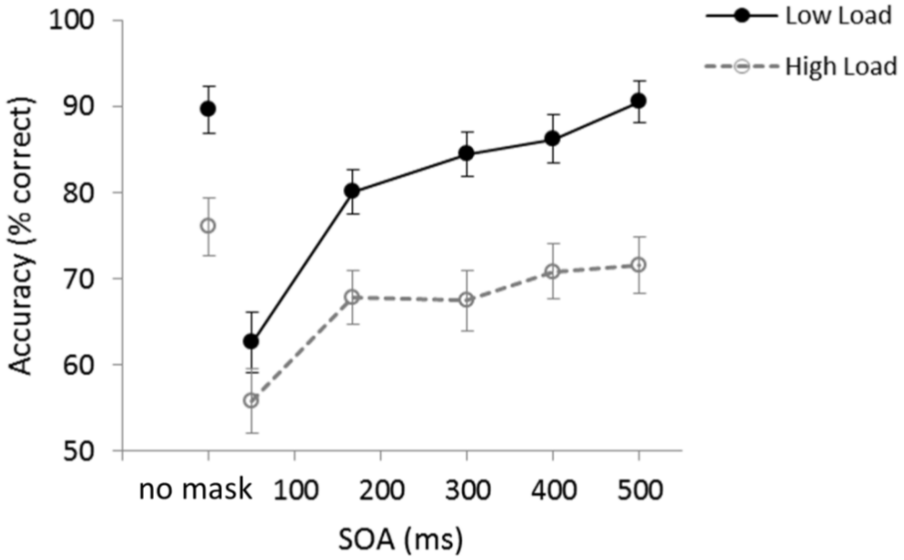

Figure 5.3. Discrimination accuracy for the masked target under low and high perceptual load in the letter search task in Experiment 11. Data were averaged over 14 participants and error bars represent ±1 *SEM*.

5.3.3 Discussion

First of all, the results replicate the main findings of the previous experiment: Overall, masking is more effective and recovery from masking is slower under high compared to low load. Secondly, performance under low load did not fully recover by 167ms SOA (did not reach an asymptote) but significantly improved until 500ms SOA. Finally, performance under high load remained suppressed even at long SOAs of 400ms or 500ms.

5.4 Experiment 12

In both previous Experiments (10 and 11), metacontrast masking was maximal at the 50ms SOA, independent of the level of load. However, the largest difference in masking strength between low and high load was always observed at later SOAs in both experiments.

This finding suggests that the effect of load in modulating metacontrast masking may be more pronounced at longer SOAs (e.g. 400ms or 500ms), due to the delayed recovery under high compared to low load. This effect at SOAs beyond 150ms indicates the involvement of higher level attentional mechanisms or recurrent feedback, formalized in the theories of *interruption masking* (e.g. Michaels & Turvey, 1979) and *object substitution* (Enns & Di Lollo, 1997) of metacontrast masking. These theories maintain that masking can occur at longer SOAs irrespective of the physical properties of the mask (which should only affect sensory interactions between target and mask that act over short distances and are therefore observed at short SOAs) Furthermore, the effect either depends on interruption of attentional mechanisms transferring information from visual iconic memory to non-visual representations (e.g. working memory) (Michaels & Turvey, 1979; Neisser, 1967; Sperling, 1960) or the early substitution of the target features by that of the mask due to rapid recurrent feedback (DiLollo et al., 2000; Enns & DiLollo, 1997). Thus, if modulation of the masking curve by perceptual load is indeed most effective at later SOAs, physically altering the metacontrast mask (such as reducing its contrast) should mainly reduce overall masking strength at early SOAs (e.g. at 50ms SOA), while leaving load effects at later SOAs (e.g. 167ms to 500ms) unaffected.

Additionally, some of the observers reached floor performance at the 50ms SOA under high load in the previous experiments. It is therefore possible that floor performance under high load may have contributed to the significant interactions found between load and SOA. To test if this was the case and to examine whether modulation of masking strength by load follows the predictions of *interruption masking* or *object substitution*, the contrast of the mask was reduced to roughly 20% of its original contrast.

5.4.1 Method

Observers. Fourteen observers (12 females, 19 to 33 years of age, *Median* = 23.5), none of which had participated in the previous experiment, participated in Experiment 12. All observers were recruited through the University College London subject pool and were naïve to the purpose of the experiment.

Stimuli and Procedure. All aspects of the apparatus, stimuli and procedure were identical to that of Experiment 10, except for the contrast of the mask which was reduced to roughly 20% of its original contrast. The SOAs used were also the same as in Experiment 10: 50ms, 167ms, 317ms, 500ms and a no mask condition.

5.4.2 Results

Visual search. As previously, under high perceptual load accuracy (M = 79%, SEM = 3.15) was significantly reduced and reaction time (M = 1121ms, SEM = 78) significantly increased compared to accuracy (M = 97%, SEM=.74) and reaction time (M = 848, SEM = 72) under low load; $t(13) = 6.17$, $p < .001$ and $t(13) = 6.55$, $p < .001$ for accuracy scores and reaction times, respectively.

Shape discrimination. The presence of the metacontrast mask significantly reduced target discrimination accuracy (M = 83%, SEM = 2.2) compared to the no mask condition (M = 86%, SEM = 1.8), confirming the efficacy of the mask irrespective of the level of perceptual load in the letter search task; *F*(1,13) = 5.52, $p < .05$, $\eta_p^2 = .30$. The same effect was observed for the level of load which significantly reduced performance for the shape discrimination under high (*M* = 76%, *SEM* = 3.4) compared to low load (*M* = 93%, *SEM* = 1.2), irrespective of the presence of the mask; *F*(1,13) = 23.27, $p < .001$, $\eta_p^2 = .64$.

Figure 5.4

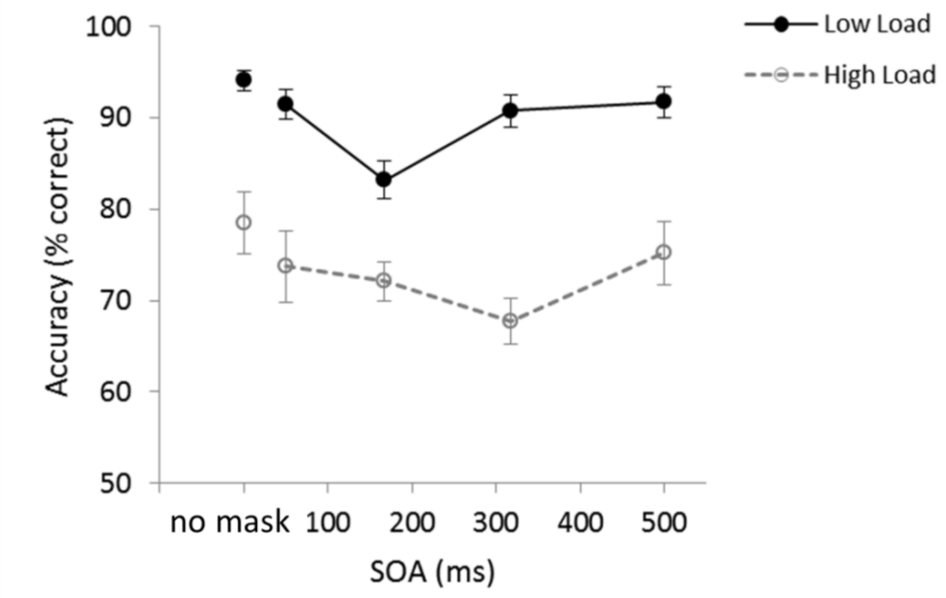

Figure 5.4. Discrimination accuracy for the masked target under low and high perceptual load in the letter search task in Experiment 12. Data were averaged over 14 participants and error bars represent ±1 *SEM.*

With reduced mask contrast, the obtained masking curves under low and high load (Figure 5.4) showed a markedly different pattern compared to the previous experiment (Experiment 10, Figure 5.2). A two-factor, within subject repeated measures ANOVA with the factors Load and SOA revealed significant main effects of

Load (*F*(1,13) = 39.18, $p < .001$, $\eta_p^2 = .75$) and SOA (*F*(4,52) = 11.91, $p < .001$, $\eta_p^2 = .48$), as well as, a significant interaction between Load and SOA; *F*(4,52) = 4.03, $p < .01$, $\eta_p^2 = .24$. Furthermore, pairwise comparisons (Bonferroni adjusted for multiple comparisons) showed that while under low load performance recovered from 167ms SOA onwards (167ms *M* = 83%, *SEM* = 2.0; 317ms *M* = 91%, *SEM* = 1.8), *t*(13) = 6.01, $p < .001$, this was not the case under high perceptual load (167ms *M* = 72%, *SEM* = 2.2; 317ms *M* = 68%, *SEM* = 2.5); *t*(13) = 1.85, *p* = .87, where performance appeared to deteriorate further until 317ms SOA. Lastly, under low load, performance at 317ms SOA (*M* = 91%, *SEM* = 1.8) had already fully recovered to the level of the no mask condition (*M* = 94%, *SEM* = 1.1), *t*(13) = 2.26, *p* = .41), while under high load recovery was only complete at 500ms SOA (*M* = 75%, *SEM* = 3.4; no mask *M* = 79%, *SEM* = 3.4); *t*(13) = 1.84, *p* = .89.

5.4.3 Discussion

First of all, the results indicate that when mask contrast is significantly reduced, the strong masking effect observed at 50ms SOA in the two previous experiments (10 and 11) disappears. Figure 5.4 illustrates how for both, low and high perceptual load, masking strength further increases at SOAs greater than 50ms and only starts to recover at around 167ms SOA in the low load and 317ms SOA in the high load condition. Secondly, the significant interaction between level of load and SOA suggests that the effect does not rely on the floor performance at 50ms SOA observed for a few observers in the previous experiments. Thirdly, performance recovered fully in both load conditions. This finding rules out any effects of response interference at the longest SOA (500ms) under high load, which could have accounted for reduced performance at this SOA in the previous experiments. Finally, compared to the respective no mask conditions under low and high load, masking was most effective at early SOAs under low load, whereas under high load masking was more pronounced at later SOAs.

Together, the results indicate that limiting attentional resources through high perceptual load in an unrelated task may lead to *interruption* or *object substitution* involving later stages of visual processing.

5.5 Experiment 13

The previous experiments showed that metacontrast masking, defined by the proximity of target and mask contours, was increased in strength and could occur at longer SOAs under high perceptual load, compared to low perceptual load in the letter search task. The findings that under high load, masking persisted at longer SOAs and its maximum effect (the lowest point in the masking curve) shifted towards longer SOAs when mask contrast was reduced, indicate that it is likely that other factors instead of, or in addition to, rapid, local, inhibitory processes in early visual cortex, contributed to the results. If this is the case, other forms of masking that are vulnerable to top-down influences should produce qualitatively similar effects. One other form of masking that is thought to strictly depend on drawing the attentional focus away from the target is *object substitution* masking (see Chapter Introduction). In this type of masking the mask is clearly separated from the location of the target and, therefore, any observed masking effects presumably do not depend on local inhibition. This type of masking is strongest when the mask appears together with the target but remains on screen after the target has disappeared (*common onset masking*; Di Lollo et al., 2000). The purpose of the following experiment was to test the effects of perceptual load on *object substitution* masking by using a mask that was clearly separated from the target (in this case, four small dots surrounding the target location).

However, all other parameters such as SOAs and mask duration were identical to the previous metacontrast experiments, such that the mask always appeared after target onset and remained on screen for the same duration as the target (33ms). This allowed for a direct comparison between the effects of load on *metacontrast* and *object substitution* masking without changing the temporal characteristics of the mask (as in *common onset masking*; Di Lollo et al., 2000).

5.5.1 Method

Observers. For Experiment 13, fourteen new participants were recruited from the University College London subject pool (8 females, 18 to 31 years of age, *Median* = 20). All had normal or corrected to normal visual acuity and were naïve to the purpose of the experiment.

Stimuli and Procedure. The stimuli and procedure were identical to the previous metacontrast experiments (Experiments 10-12) except that the mask now only consisted of four small white dots (0.5° in diameter) surrounding the target at 1.3° eccentricity (Figure 5.5). The SOAs between target and mask were 50ms, 167ms, 300ms, 400ms, and 500ms; plus a no mask condition.

Figure 5.5

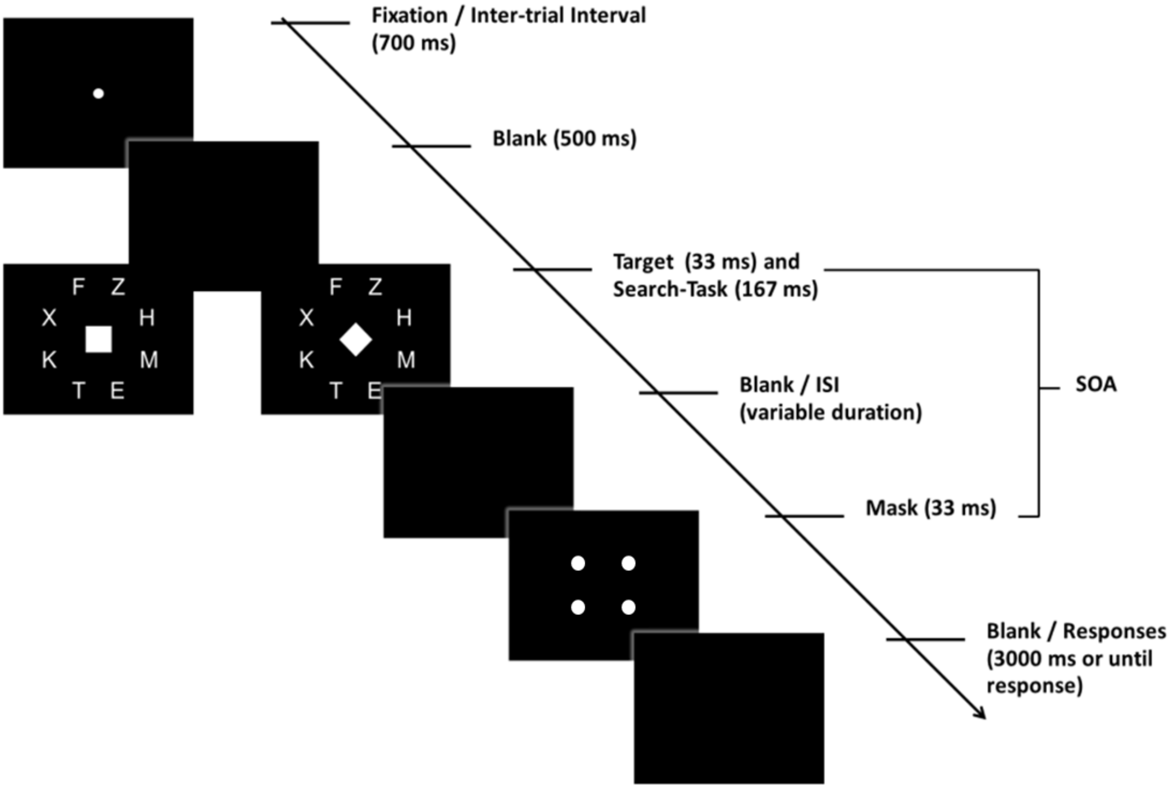

Figure 5.5. Trial sequence in Experiment 13. Participants performed a letter search task (high load shown here) while also discriminating the central target (diamond vs. square). The mask consisted of four small, white dots, forming a square around the centre. The mask was clearly separated from the location of the target (by ~1°) and appeared between 50 and 500ms after target onset.

5.5.2 Results

Visual search. Performance in the letter search task significantly decreased under high vs. low load, with lower accuracy (*M* = 78%, *SEM* = 1.83) and longer reaction time (*M* = 1328ms, *SEM* = 76) under high load compared to low load (accuracy *M* = 97%, *SEM* = .72; reaction time *M* = 1034, *SEM* = 71); *t*(13) = 10.67, *p* < .001 and *t*(13) = 6.24, *p* < .001, for accuracy and reaction time respectively.

Shape discrimination. The 4-dot mask was effective in significantly reducing discrimination accuracy (*M* = 83%, *SEM* = 2.2) compared to when no mask was presented (*M* = 92%, *SEM* = 2.0); *F*(1,13) = 41.46, *p* < .001, η_p^2 = .76. Furthermore, high perceptual load in the letter search task reduced target discriminability overall (*M* = 81%, *SEM* = 2.7) compared to low load (*M* = 93%, *SEM* = 1.6), irrespective of the presence of the mask; *F*(1,13) = 32.58, *p* < .001, η_p^2 =.72.

Figure 5.6 illustrates the effects of low and high perceptual load on the masking function obtained with the four-dot mask.

Figure 5.6

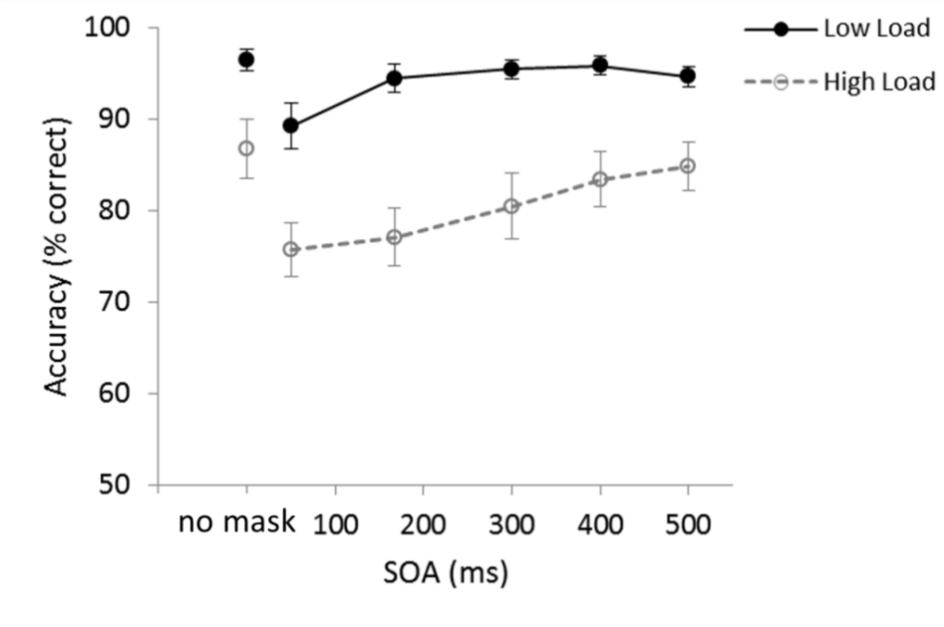

Figure 5.6. Discrimination accuracy for the masked target (object substitution masking) under low and high perceptual load in the letter search task in Experiment 13. Data were averaged over 14 participants and error bars represent ±1 *SEM*.

The results show minimal masking under low perceptual load confined to the 50ms SOA. Under high load, however, depression of target discrimination accuracy was still evident at longer SOAs. A two-factor, within subject repeated measures analysis of variance (ANOVA) on target discrimination accuracy with the factors Load and SOA revealed significant main effects of Load, $F(1,13) = 43.07$, $p < .001$, $\eta_p{}^2 = .77$, and SOA, $F(2.46,31.98) = 10.39$, Greenhouse-Geisser corrected, $p < .001$, $\eta_p{}^2 = .44$, as well as, a significant interaction; $F(5,65) = 2.39$, $p < .05$, $\eta_p{}^2 = .15$.

Furthermore, pairwise comparisons (with Bonferroni adjustment for multiple comparisons) confirmed that under low load accuracy fully recovered to the level of the no mask condition at 167ms SOA (no mask *M* = 97%, *SEM* = 1.1; 167ms *M* = 95%, *SEM* = 1.5), $t(13) = 1.82$, $p = 1.00$; while this was not the case under high load (no mask *M* = 87%, *SEM* = 3.1; 167ms *M* = 77%, *SEM* = 3.1), $t(13) = 3.83$, $p < .05$.

5.5.3 Discussion

Together the results of Experiment 13 demonstrate that perceptual load significantly modulates the effectiveness of a 4-dot mask, indicating that even when mask and target are spatially separated, high perceptual load increases masking strength. Moreover, the results are in line with the previous experiments, indicating not only increased masking strength under high load but also slower recovery to the level of the no mask condition. As the results were obtained using a mask that does not rely on spatial proximity to the target or shared low level visual features between target and mask, the findings suggest that perceptual load also modulates target-mask interactions at later stages in the visual system.

5.6 Chapter Conclusions

Traditionally, metacontrast masking has been conceptualized as a purely feedforward process involving inhibitory interactions among local contours (e.g. Breitmeyer, 1984). More recently, however, this view has been challenged by studies demonstrating attentional effects on the strength of metacontrast masking (e.g. Di Lollo et al., 2000). Many authors have taken these findings to imply that metacontrast masking cannot depend solely on local processes in low level vision but must also occur at higher levels of visual processing. Others have argued that although attention is likely to play a role, it may do so by simply influencing the proposed mechanism of contour interactions in early vision (e.g. Breitmeyer & Ogmen, 2006). In any case, the numerous findings of attentional effects warrant an inclusion of attention within any complete theory of metacontrast masking.

The results from the current chapter extend this research by demonstrating modulations of metacontrast masking that depend on the level of perceptual load (and thus, restriction of attentional resources) in an unrelated task. This finding is important with regard to previous studies investigating attentional effects on metacontrast masking, as it suggests that perceptual integration of target and mask can be modulated without modulating spatial attention or set size. Moreover, the large majority of previous studies involving attentional modulations, varied display complexity (e.g. pop-out, set size) while the target was always part of the display. In these studies the target was only identifiable after presentation of a spatial cue to the target location (e.g. Tata, 2002) or retrospectively after presentation of the mask (which appeared only at the location of the target) (e.g. Spencer & Shuntich, 1970). In the experiments reported here, however, attention was modulated in a completely unrelated task while keeping all target and mask stimulus parameters constant. In contrast to the results of the previous chapter (pattern masking), with metacontrast masking the difference in masking strength between low vs. high perceptual load was most pronounced at later SOAs (greater than 100ms). This effect is akin to the proposal of Enns & DiLollo (1997), who described the particularly strong masking that occurred when attention was directed away from the masking stimulus (*object substitution)*. The authors attributed the masking effect to higher level visual processes involved in object recognition and dissociated it from low level mechanisms by using a mask consisting of only 4 small dots surrounding the target. With this type of mask, when attention was directed away from the target location, the effect was comparable to that obtained with a standard metacontrast mask without manipulating attention. However, in studies using the 4-dot mask, attention was always manipulated by increasing the set size (number of possible target locations), often comparing a large display to a single item condition. Thus, the authors compared focussed attention on a single item with distributed attention to many items where the target was only identifiable *after* the presentation of the mask.

To probe the effects of perceptual load on *object substitution* masking, the last experiment in the current chapter used the identical design as the previous experiments but exchanged the metacontrast mask with 4 small dots. This design provides a measure of *object substitution* masking (and the dependency of metacontrast masking on contour interactions) under high and low perceptual load while maintaining identical stimulus parameters for target and mask across the different attention conditions. The results show that perceptual load indeed increases masking with 4 dots, demonstrating a slower recovery to the level of the no mask condition under high compared to low load.

Together, the results suggest that the effectiveness of a metacontrast mask is increased and maintained throughout longer SOAs by high perceptual load in an unrelated task. Furthermore, the findings demonstrate that the effectiveness of a mask that is spatially separated from the location of the target (4-dot mask) can be modulated by attention (here, perceptual load) even when target and mask are always presented at the same spatial location, and target onset timing, as well as, mask duration are constant throughout the experiment.

Chapter 6

General Discussion

6.1 Summary of the findings

The primary aim of this thesis was to elucidate some of the underlying mechanisms leading to reduced perception under conditions of high perceptual load for basic visual processes. Specifically, the influence of visual context on target perception was assessed under different load conditions, hypothesising that perceptual load should not only affect target stimulus processing alone, but also the interaction with other stimuli (which can be conceived of as external noise).

Using a well-established noise masking paradigm (e.g., Baldassi & Verghese, 2005; Blake & Holopigian, 1985; Legge & Foley, 1980; Ling & Blake, 2009; Ling et al., 2008; Majaj et al., 2002; Solomon & Pelli, 1994) to measure orientation tuning psychophysically (in Chapter 2), high perceptual load was shown to reduce the overall signal strength of orientation responses across the tuning profile. However, in addition to a reduction in orientation signal strength, the precision of the signal was reduced, due to broader orientation tuning under high compared to low load. This is akin to a greater impact of noise orientations in the stimulus representation under high load. The finding also generalized to a different visual search task where the orientation content was equal in the two load conditions and also did not depend on reductions in the apparent contrast of stimuli under high load.

Furthermore, examining the effect of load on contextual effects that purely rely on the interactions among orientations (for example, the tilt-illusion) confirmed reduced noise suppression under high load and the dependence of contextual modulation on stimulus configuration (Chapter 3). Perceptual load was found to alter the appearance of the tilt-illusion, reducing the magnitude of the illusion under high compared to low load. This finding indicates that under the right conditions, limiting perceptual resources may increase veridical perception by reducing the strength of illusions that rely on the same resources. However, the effect of load strictly depended on the stimulus configuration and conscious perception of the context: When target and inducer orientations in the tilt-illusion spatially overlapped, high perceptual load reduced the magnitude of the illusion. However, when target and inducer were spatially separated and the inducer was rendered “invisible” by backward masking, the effect of load was reversed, increasing the magnitude of the illusion under high compared to low load. This effect did not depend on the physical properties of the stimuli (e.g. contrast, spatial frequency, noisiness) and was only observed when the inducer orientation was not consciously perceived and target and inducer orientations were spatially separated (Chapter 3). Together these findings suggest an alternative account for the effects of perceptual load on contextual integration in terms of the increased impact of task irrelevant stimulus components (the context) under high load (discussed in more detail below).

Moreover, a reduction in the ability to separate signal from noise under high load was also found (Chapter 4 and 5) when noise was introduced after presentation of the target (i.e. in backward masking) instead of presenting target and context simultaneously. Here the effect of load on the time course of contextual integration was examined under several stimulus conditions: Target and mask either shared similar features and were presented in the same spatial location (Chapter 4), shared some features but were presented in adjacent locations or did not share any features and were presented in spatially separate locations (Chapter 5). The results demonstrated that high perceptual load affected the interaction between target and mask by increasing the overall effectiveness of the mask but also delayed recovery from masking. The pattern of results suggested an early locus for perceptual load effects at the stage of low level signal accumulation but also showed later effects of load, suggesting modulation at higher levels in the visual system involved in object recognition.

6.2 Implications for perceptual load theory

The findings merit an extension of perceptual load theory to incorporate the influence of the visual context on perception of unattended stimuli. The reported reduction in signal strength under high load is in line with numerous previous reports of reduced perceptual sensitivity (e.g. Cartwright-Finch & Lavie, 2007; Macdonald & Lavie, 2008), reduced adaptation (Bahrami et al., 2008), reduced neural excitability (Muggleton et al., 2008) and reduced stimulus evoked activity (e.g. Bahrami et al., 2007; Rees, et al., 1997; Schwartz et al., 2005) under high load.

However, the additional effect of broadened tuning under high load provides an alternative explanation for the reduced BOLD responses obtained with fMRI. With increased noise due to imprecise tuning in conditions of inattention under load, visual cortex response can no longer be discriminated from baseline levels of activity. The result therefore provides a new, hitherto undiscovered mechanism for perceptual load effects which previous studies failed to detect due to inherent limitations of behavioural measurements or the limited spatial resolution of neuroimaging techniques. Furthermore, the perceptual consequences of broadened orientation tuning in combination with reduced gain are likely to underlie many of the perceptual phenomena associated with reduced perception under high load, such as *inattentional blindness* (e.g. Cartwright-Finch & Lavie, 2007; Macdonald & Lavie, 2008) or *change blindness* (Lavie, 2006).

The observed changes in the width of psychophysical orientation tuning curves are likely the result of modulation in lateral cortical interactions, thought to shape the population response to orientation (e.g. Blakemore & Tobin, 1972). While such interactions do not affect the orientation tuning bandwidth of individual neurons (e.g. Ferster & Miller, 2000) they have been shown to extend over relatively large cortical distances (e.g. Joo, Boynton & Murray, 2012) and are thought to underlie contextual effects involving interactions among orientations (Blakemore & Tobin, 1972). Thus, if perceptual load alters large scale cortical interactions among orientation detectors, this should be evident in modulations of contextual orientation integration. The reported effects on the tilt-illusion (a paradigmatic example of such contextual effects) where orientations appear shifted away from the orientation of the context (tilt-repulsion), confirm this hypothesis. These results further extend the scope of perceptual load effects by demonstrating changes in low level interactions between target and context.

The observed changes to spatial interactions due to reduced noise suppression suggest perceptual load may also influence the ability to separate signal from noise even when they do not occur simultaneously. Modulation of contextual effects by changes in lateral inhibition in early visual cortex, for example, has been shown to be strongly influenced by top-down recurrent feedback (e.g. Gilbert & Li, 2013). It is therefore important to consider perception as a dynamic process that evolves in the brain over time (Van Rullen & Thorpe, 2001), mediating task dependent, attention driven modulations. Therefore, any effects of perceptual load on orientation responses must have a temporal signature that can inform our understanding of the visual stages involved in mediating the effects on perception. The time course of orientation detection and discrimination under load was examined by using an orientation band-passed noise mask presented at various SOAs after target presentation. The results revealed reduced but also slower target processing under high compared to low load. Interestingly, this was the case for both orientation detection and discrimination. This finding indicates a relatively early locus for the effects of high perceptual load in pattern masking, as the observed effects of load did not shift or extend to longer SOAs in orientation discrimination compared to detection. As target detection is a necessary first step for discriminating its orientation, load appears to modulate this early stage of visual processing in both tasks. However, further testing showed that perceptual load also modulates the effectiveness of masks that do not spatially overlap the target (*metacontrast masking*) and is still effective when mask and target are well separated and do not share contour orientations (*object substitution masking*). These findings suggest interruption of target processing at higher levels in the visual system.

Together, the findings reported in this thesis suggest that the profound changes to perception observed under perceptual load are at least in part mediated by changes in low level vision, specifically the neural selectivity of orientation responses. These effects have been tentatively localized to the population response in early visual areas, by demonstrating changes in contextual interaction under load and demonstrating prolonged accumulation of low level visual information under high compared to low load.

6.2.1 Neurophysiological correlates of perceptual load effects

What could be the cause of the broadened tuning under high perceptual load? One possible explanation is that perceptual load alters the tuning width of individual neurons. Originally, orientation tuning at the single neuron level was thought to arise purely from the physical arrangement of connections between neurons in lateral geniculate nucleus and cells in V1 (Hubel & Wiesel, 1962; Shou & Leventhal, 1989). Such a simple feedforward model, however, cannot explain many of the empirical observations made regarding orientation selectivity. For example, it cannot account for the invariance of orientation tuning with increased stimulus contrast (Sclar & Freeman, 1982). Due to the discrepancy with empirical findings, some have proposed modified versions of the original feedforward model (Ferster & Miller, 2000; Priebe & Ferster, 2008) that incorporate non-linear properties of feedforward neural responses required to further sharpen the broad tuning biases from thalamocortical inputs. Others, however, have argued that precise orientation tuning requires lateral inhibition in cortical visual areas or feedback (Marino et al., 2005; Shapley, Hawken & Ringach, 2003; Somers, Nelson & Sur, 1995; Sompolinsky & Shapley, 1997). However, although attentional effects have been reported in areas even as early as the LGN (O'Connor et al., 2002), no single neurophysiological study has reported convincing evidence for a change in orientation tuning bandwidth at the level of individual cortical neurons, whereas many have reported a change in gain without a change in tuning for both, spatial and feature based attention (e.g. Crist et al., 2001; Maunsell & Treue, 2006; McAdams & Maunsell, 1999). Therefore, the combined response at the level of the population is more likely to determine perceptual consequences of attentional manipulation (including manipulations of perceptual load reported here).

In what ways could perceptual load affect the population response to unattended stimuli? As discussed previously, feature-based attention to motion direction has been shown to result in sharper tuning at the neural population

(Martinez-Trujillo & Treue, 2004). Although the decisive factor influencing behaviour is the sharpened tuning found at the population level, this and similar findings have been interpreted in terms of the classical approach to attentional improvements in behaviour, that is the result of an improved signal-to-noise ratio by increasing the firing rate of individual neurons (e.g. the *feature-similarity gain model*; Martinez-Trujillo & Treue, 2004; Treue & Martinez-Trujillo, 1999). Here, sharpened tuning at the population is explained by modulation of the gain in individual neurons depending on their feature preferences. More recently, however, research has shifted toward understanding the interactions among neurons that cause changes at the population level which extend beyond the expected benefits from changes in firing rate found in single neurons. This approach has demonstrated, for example, that spatial attention modulates the correlation in firing rate between pairs of neurons (Cohen & Maunsell, 2009). The read out of information from populations of neurons depends not only on the variability of responses from single neurons but also on how much the variability of individual responses correlates among groups of neurons. If the noise in single neurons (expressed as their variance in firing rate) would be independent, it would be completely averaged out at the level of the population. However, as trial-by-trial noise correlations between pairs of neurons do occur (Shadlen, Britten, Newsome & Movshon, 1996; Shadlen & Newsome, 1998; Zohary, Shadlen & Newsome, 1994), any decrease in such correlations should result in a better read out of the population mean. For a given attended feature (e.g. a vertical orientation) this would be represented by a sharper psychophysical tuning curve. Indeed, it has been shown that decreases in noise correlations due to spatial attention can account for up to 80% of the improvement in the population signal, while increases in firing rate of individual cells account only for a much smaller fraction (Cohen & Maunsell, 2009).

These findings hint at a much larger role of neural interaction in shaping the population response mediated by attention than was previously assumed, suggesting an increase in noise correlations between pairs of neurons may represent a plausible explanation for broadened orientation tuning under high perceptual load.

6.3 Implications in relation to previous research on attention

6.3.1 Psychophysical investigations of gain and tuning

6.3.1.1 Spatial cueing

A number of previous studies have investigated the effects of other attentional manipulations, such as spatial and feature based attention, on psychophysical tuning curves. Recall that attention may modulate the gain of a tuning curve, increasing the strength of the signal, or change its tuning width which determines the suppression of irrelevant noise. In contrast to the current findings, in these studies, spatial attention consistently increased the gain, multiplying responses by a constant factor across the tuning curve, without modulating tuning bandwidth. This was shown for motion perception (Ling et al., 2009) using indirect measures of tuning by comparing responses under low and high motion coherence (equivalent to high and low external noise), as well as for orientation, using a more direct measure similar to the noise masking technique reported in Chapter 2, where noise is presented together with the signal, impairing its detection or discrimination (Baldassi & Verghese, 2005). Both of these studies (Baldassi & Verghese, 2005; Ling et al., 2009) concluded that perceptual benefits of spatial attention are mediated by increasing the gain across the population and not by changes in tuning. Consistent with this interpretation, other studies have shown that perceptual tuning is unaffected by spatial attention. These have used reverse correlation techniques to estimate the properties of the perceptual filter used by the observer (*noise image classification*; Eckstein et al., 2002; Murray et al., 2003; Neri, 2004) or changed the bandwidth of the noise to estimate how this changes threshold elevation (*critical-band masking*; Lu & Dosher, 2004; Talgar, Pelli & Carrasco, 2004).

First of all, the lack of observed changes in tuning with previous manipulations of spatial attention compared to the significantly broadened tuning under high load reported here, suggests that the observed changes must result from the level of load, and presumably the restriction of perceptual resources, and are not due to different allocations of spatial attention in the two load conditions. Secondly, the current results suggest that perceptual load is a necessary requirement for inducing changes in sensory tuning, whereas spatial cueing appears neither necessary nor sufficient.

6.3.1.2 Feature cueing

For the effects of feature based attention on tuning curves, evidence for both, a change in gain and tuning has been reported and no clear consensus has been reached within the field. For example, two of the studies discussed above (Baldassi & Verghese, 2005; Ling et al., 2009) also estimated the effect of feature based attention on tuning curves. Both concluded that feature based attention was mediated through a change in tuning. Ling and colleagues (2009), however, found an increase in gain in addition to sharper tuning for feature based attention. Others have found only an increase in gain, but no change in tuning when feature based attention was manipulated by attending to a set of moving dots in the presence of distractor dots (Murray et al., 2003). Firstly, these findings support the notion that attentional factors can influence the precision of the population response to orientation (or motion) by changing the width of psychophysical tuning curves. Secondly, the effects on tuning width reported previously are limited to feature based attention. In these studies feature based attention is assessed by comparing performance in conditions where a specific feature (e.g. a specific orientation) is either validly cued before it appears or is invalidly cued (or not cued at all). Thus, the studies rely on manipulating the expectation (or top-down task set) for a specific feature to be presented. This manipulation is different from the dual-task paradigm used throughout the experiments in this thesis. Here, no cues were used to manipulate spatial or feature based attention. Moreover target location and target features were always identical on each trial throughout a given experiment. Recall from Chapter 2, for example, that the orientation target was always vertical and its feature relevance was the same in the two load conditions. Thus, the instructions for target detection were equal and the top-down relevance for the feature target was unaffected by the level of load.

6.3.1.3 Dual-task paradigms

In order to clarify the role of perceptual load in tuning, it is important to compare the results to other studies that also employed dual task paradigms to investigate changes in tuning. Busse and colleagues (Busse, Katzner, Tillmann & Treue, 2008), for example, used a dual task paradigm where participants detected large (*full attention* condition) or small (*poor attention* condition) changes in luminance at fixation, while responding to a predefined target direction of moving dots which was presented occasionally within a series of brief, fully coherent motion pulses in other, non-target directions. The results showed a multiplicative increase in gain (when the motion direction was fully attended), but no change in direction tuning width of the obtained tuning curves. The discrepancy between these results and those reported in this thesis may be due to the different stimulus dimension examined (motion direction vs. orientation) but could also be due to the use of an inefficient load manipulation that did not fully engage attentional resources. This claim is supported by examining the reported accuracy scores for the luminance detection task at fixation. Performance for this task in their *poor attention* condition did not significantly differ from performance in the *full attention* condition (88.5% and 85% correct responses in the poor and full attention conditions respectively). Thus, the efficacy of the attentional manipulation in critically limiting attentional resources was not warranted.

Lee et al. (1999) conducted a similar study in the domain of orientation perception. Here, participants performed a demanding visual search task at fixation while detecting an oriented target embedded in a noise mask. Orientation tuning curves obtained in the dual-task condition showed reduced gain and broader tuning compared to those obtained from a single task condition where attention was not focused on the visual search task. These findings are consistent with the results reported in this thesis. However, Lee and colleagues (1999) compared single and dual-task conditions. This comparison could yield changes in tuning simply due to the additional processing of orientations (for the visual search task) which would not reflect changes due to attentional modulation. Moreover, the dual-task may have resulted in additional demands on memory and response compared to the single-task condition.

In summary, the reported changes in orientation tuning under high perceptual load challenge previous research that concluded attention has no effect when cued to locations (or, in some studies, to features) and is purely mediated by a change in gain. As the change in tuning reported here was obtained in the absence of any spatial or feature based manipulations and the target was always presented in the same spatial location, having the same orientation (vertical), it is possible that the observed changes in population tuning require a certain level of perceptual load in a separate task. This suggests that while perceptual load appears to be a necessary condition for the effect of attention on tuning, spatial cueing is neither necessary nor sufficient for an effect on tuning.

6.3.2 Attentional effects on contextual integration and cortical interaction

As detailed previously, the population response to orientation is determined by lateral interactions between orientation columns in early visual cortex (Blakemore & Tobin, 1972). Neurons with different orientation preferences ordinarily inhibit one another, shaping the tuning profile at the population level.

The effects of perceptual load on contextual integration reported in Chapter 3 are consonant with an interpretation of load modulating the balanced inhibition at the level of the neural population. In agreement with broadened tuning due to reduced lateral inhibition among orientation detectors, high perceptual load in a separate task resulted in a reduction of the perceived tilt-illusion. However, this was only found when there was a specific requirement for distractor exclusion, for example, when target and inducer were presented superimposed within the same spatial location. In this case the target orientation had to be actively distinguished from the inducer orientation in order to identify its orientation with respect to subjective vertical. When target and inducer were spatially separated, however, they could be identified as such simply based on their spatial location, reducing the draw on additional resources required for target orientation detection. This interpretation suggests that contextual modulation is stronger when the context is at least partly task relevant or must be distinguished from the target in order to perform the task successfully. Other studies, in support of this hypothesis, have found that attentional effects on contextual integration depend on stimulus configuration and the specific task performed. For example, attentional effects increase when multiple stimuli compete (e.g. Kastner et al., 1998; Motter, 1993) or stimulus complexity is increased (Motter, 1993) and depend on the geometric relationships among stimulus features (e.g. colinearity; Li, Piech & Gilbert, 2006). In addition, it has been shown that the tuning properties of cells in monkey V1 can change depending on task relevance of different stimulus components, even when the same stimuli are used (Li, Piech & Gilbert, 2004).

However, an alternative explanation of the findings of Chapter 3 must be considered based on the reduced ability to separate the signal from noise under high perceptual load (observed in Chapter 2). When target and inducer were spatially overlapping, observers had to identify the target first (i.e. separate it from the noise inducer) before a judgment on its orientation could be made. Although the distinction between target and inducer was aided by different contrasts and spatial frequencies, the high perceptual load task may have caused observers to erroneously respond to the inducer instead of the target orientation when the perceptual separation of target and inducer failed. This would have resulted in an artificial reduction in the tilt-illusion under high load that would not represent a reduction in the perceived magnitude of the illusion. If this was the case, perceptual load only interacted with contextual integration of target and inducer when the inducer was spatially separated and weakened by the subsequent presentation of the noise mask. In this case observers did not have to actively distinguish the target and inducer orientations as they were unaware of the presence of the inducer or its orientation. Moreover, the observed effect showed a reversal of the previous result by indicating a significant increase in the perceived magnitude of the illusion. An explanation for this result comes from previous studies that have found that distractors presented below a certain intensity threshold can exert a stronger effect on perception than suprathreshold stimuli (Tsushima et al., 2006; Tsushima et al., 2008). This finding has been interpreted in terms of reduced active suppression of “unnoticed” distractors, indicating that the attention system requires a minimum stimulus intensity for successful distractor suppression. Limiting attentional resources under high load (in Chapter 3) may have either reduced the impact of the mask (allowing stronger unconscious modulation by the inducer) or reduced active suppression of the subthreshold inducer, resulting in a stronger illusion under high load.

6.3.3 Attentional effects on the time course of spatial integration

Chapters 4 and 5 demonstrated that perceptual load influences the time course of perceptual processes involved in target detection and discrimination. This suggests that the reduced ability to separate signal from noise under high load and the concomitant increased impact of noise and visual context on target perception extends to cases where signal and noise are presented successively.

Attentional effects on masking curves have been assessed in a number of previous studies. The vast majority of these studies (discussed in the general introduction; e.g. Averbach & Coriell, 1961; Boyer & Ro, 2007; Di Lollo et al., 2000; Enns & Di Lollo, 1997; Michaels & Turvey, 1979; Ramachandran & Cobb, 1995; Shelley-Tremblay & Mack, 1999; Spencer, 1969; Spencer & Shuntich, 1970; Tata, 2002; Tata & Giaschi, 2004), however, utilized perceptual grouping of display elements or changes in display complexity to assess the effects of attention. Due to the stimulus configurations used in these experiments it is difficult to pinpoint the attentional mechanisms underlying the observed effects on masking curves.

Ramachandran & Cobb (1995), for example instructed observers to perceptually group different sets of stimuli presented simultaneously. The effect of metacontrast masking was reduced when observers perceptually grouped a masked target and a non-target but not when they grouped two non-targets. The authors interpreted the result in terms of object based attention, suggesting that masking was only reduced when the target was part of the perceptual group, due to the attentional benefit associated with feature binding. However, an explanation in terms of spatial attention is equally likely, as grouping the non-targets moved the spatial centre of the group away from the mask, while grouping the target with one of the non-targets shifted the centre of spatial attention closer towards the location of the mask (Breitmeyer & Ogmen, 2006). Other studies manipulated attention by either increasing the number of target items in the display (Averbach & Coriell, 1961; Di Lollo et al., 2000; Spencer, 1969; Spencer & Shuntich, 1970; Tata, 2002; Tata & Giaschi, 2004) or increasing the number of possible locations where the target-mask sequence could appear (e.g. Enns & Di Lollo, 1997). Under such conditions, the behavioural benefits attributed to attention can be well explained in terms of spatial or decision uncertainty (Baldassi & Burr, 2000; Cameron et al., 2004; Kinchla, 1992; Morgan et al., 1998; Shiu & Pashler, 1994; Solomon et al., 1997). According to these models, increasing the number of distractors (or empty target locations), increases the amount of noise that can interfere with identification of the target signal. First of all, this indicates that the general interpretation of many previous results in terms of attention is not warranted. These studies assumed that when the stimulus display consists of a single target element, attention can be focussed on the target location alone, whereas in multi-element displays attention is distributed. However, regardless of any attentional involvement, this type of stimulus manipulation inevitably introduces additional noise from the non-target locations in multi-element displays compared to the presentation of a single target. Therefore, the possible contribution of attentional effects in these studies cannot be assessed *per se*. Secondly, if attention indeed plays a role in single vs. multi-element displays the mechanism by which attention improves performance is difficult to isolate. Enns & Di Lollo (1997), for example, presented a target-mask sequence in one of three possible locations and compared performance to a condition where target and mask were presented in a single location. In such a comparison it is unclear whether the attentional benefit in the single location condition is due to the exclusion of noise from other locations (which becomes less efficient as the number of possible target locations increases) or is due to an actual enhancement of target processing. Thirdly, many of the previous studies cannot distinguish between attentional effects influencing target sensitivity directly (irrespective of the mask) and those affecting the processing of the mask. In all studies discussed previously that compared single vs. multi-element displays, the target was always part of the multi-element display prompting visual search. In addition, in many previous experiments the mask served as a cue for the target. Thus, while in the single element display the target location could be attended *before* the mask-target sequence appeared, in multi-element displays the target location was only known *after* the presentation of the mask. These considerations suggest that mask saliency and associated attentional capture (e.g. Yantis, 1993) may influence performance in addition to processes related to perception of the target. In line with this suggestion, Tata and Giaschi (2004) showed that when a mask appeared at all possible target locations in a multi-element display (with a pre-defined target) performance improved (at SOAs greater than 200ms) compared to when only a single mask was presented at the target location. The authors reasoned that the increase in mask set size eliminated attentional capture of the mask at the greater SOAs. In light of this finding it is not surprising that some types of masking where the mask remains on screen for extended periods of time after the target has disappeared (e.g. *common-onset masking*; Di Lollo et al., 2000) are particularly prone to attentional capture by the mask and thus produce stronger and more extensive suppression of target processing at later SOAs.

The masking experiments reported in Chapters 4 and 5 of this thesis extend the previous research by assessing the effects of perceptual load on backward masking, while avoiding any of the above confounds relating to spatial uncertainty, distractor noise or mask saliency. First of all, perceptual load was independently modulated in a separate task. Thus, the masked target stimulus was not part of the visual search display determining the level of load. This allowed for a direct comparison between the effects of low and high load, while keeping the number of targets and distractors in the visual search task, as well as, in the central shape detection or discrimination task identical. Furthermore, there was always only a single central target which appeared in the same spatial location with a predictable onset in each trial. Finally, the mask only appeared at a single location (that of the central target) and did not vary in duration. Thus, the reported effects on masking strength and temporal extent of masking can be directly attributed to the level of load in the search task, limiting available resources.

The findings from the pattern masking experiments in Chapter 4 suggest an early locus for perceptual load effects when target and mask are similar in structure and intensity as the strongest deviation in the pattern of results between high and low load occurred at very short SOAs, between 50ms and 100ms. Importantly, this was the case for both, orientation detection and discrimination. This suggests that perceptual load affected target processing at the stage of stimulus detection before further processing required for discrimination occurred. Furthermore, the results specifically implicate interruption of target processing as the underlying mechanism instead of degradation of the target by simple sensory integration of target and mask (i.e. *integration masking*). Together the results support the modulation of recurrent feedback (e.g. Lamme, 2000) which has been shown to affect early visual responses signalling contextual interactions (e.g. figure ground segregation) but not simple responses to orientations in isolation. The results from Chapter 5 indicated that perceptual load also modulated processing at even later stages in the visual system. Here, the effects of load were still found at SOAs beyond 300ms and by using a mask that did not rely on local contour interactions.

In all of the masking experiments high perceptual load did not only increase the effectiveness of the mask but also prolonged its effect. Thus, when perceptual resources were limited, completion of target processing required a longer period of time. By using a number of different masking techniques and SOA ranges, the results suggest that target processing is prolonged under high load throughout the visual hierarchy.

6.4 Alternative accounts of the findings

A limiting factor in all reported experiments may be the manipulation of perceptual load by modulating the heterogeneity of target and distractors in the visual search display. This particular visual search paradigm was carefully chosen as a well-established paradigm used to manipulate perceptual load that has been shown to reliably reduce distractor interference (Beck & Lavie, 2005; Forster & Lavie, 2007; 2008; Lavie, 1995; Lavie & Cox, 1997; Lavie & Fox, 2000) and sensitivity to stimuli in a secondary task (Cartwright-Finch & Lavie, 2007; Macdonald & Lavie, 2008).

Moreover, the effectiveness of the task to suppress irrelevant processing has been shown to directly depend on the efficiency of the visual search (Lavie & Cox, 1997; Roper et al., 2013). When stimuli are used that result in very efficient search (i.e. a flat search slope with increasing set size), distractor interference is enhanced compared to stimuli that produce less efficient search (i.e. a steeper search slope). Moreover, the effect of search efficiency on distractor interference has been distinguished from the effects of general task difficulty (defined as the intercept of the search function with increasing set size) which does not influence distractor interference (Roper et al., 2013).

However, it remains a possibility that some of the observed effects of perceptual load on contextual interaction depend on the physical stimulus differences between the visual search task used under low and high load. Future research should address this by employing other perceptual load tasks that use the exact same stimuli in different load conditions. In this case perceptual load can simply be determined by the task requirements within a given block of trials. For example, presenting a single, coloured cross at fixation on each trial and instructing participants to respond either to any red cross (the low load condition) or to any upright yellow or inverted green cross (the high load condition), results in the same reduced processing of irrelevant stimuli as seen with more conventional load tasks (e.g. Schwartz et al., 2005; Bahrami et al,. 2008).

6.5 Future research

Due to the invasive nature of cell recording, electrophysiological measures of neural tuning (or contextual interaction) cannot be obtained in the intact human brain. The current findings, however, suggest that conditions that constitute high perceptual load for non-human primates would also lead to attentional modulation of orientation tuning. In the human brain this can be indirectly deduced through neuroimaging research using decoding techniques to assess tuning at the neural population (rather than single cell) level (Kamatani & Tong, 2006). Specifically, estimating voxel-based orientation tuning functions (e.g. Serences, Saproo, Scolari, Ho, & Muftuler, 2009) should reveal broadened tuning under high perceptual load. Additionally, quantifying the effects of high perceptual load on the gain and tuning of neural populations in different brain areas should provide an interesting area for future research. Pratte and colleagues (Pratte, Ling, Swisher & Tong, 2013), for example, very recently showed, using neuroimaging and multivoxel pattern analysis, that in early visual areas object representations were enhanced by attention only under high levels of external noise (suggesting a noise filtering mechanism in line with the changes in tuning reported in this thesis) while in higher visual areas (V4, fusiform face area and lateral occipital complex) attention only enhanced processing under noise free conditions (indicating a change in gain without changes in tuning at higher levels of representation). Applying this methodology to research on perceptual load would yield a measure of load effects in terms of gain and tuning throughout the visual hierarchy.

Furthermore, recent computational models suggest the type of modulation at the population level appears to be relatively flexible depending on the specific requirements of the task. This flexibility could be affected by limiting available resources under high perceptual load. For instance, a number of recent theoretical models (e.g. Jazayeri & Movshon, 2006; 2007; Navalpakkam & Itti, 2007) have incorporated the finding that in visual search, boosting responses of neurons precisely tuned to the target feature is not always optimal. For example, when discriminating fine orientation differences such as distinguishing a grating tilted 5° from another grating tilted 0° (vertical), the most informative neurons are not those tuned to the 0° or 5° orientations but actually neurons tuned to roughly 15°, as these show the biggest difference in firing rate between stimuli of 0° and 5° (Regan & Beverly, 1985). It was recently confirmed that human observers can indeed employ this optimal strategy by adaptively allocating the gain to specific neural populations depending on the task (Scolari & Serences, 2009). In coarse orientation discrimination, contrast detection thresholds were lowest for the target orientation and higher for flanking distractor orientations. However, when observers performed a fine discrimination task, contrast thresholds were lower for similar, flanking distractor orientations than for the target orientation. This implies that in the coarse discrimination task observers relied on the read-out from neurons tuned to the target orientation whereas in the fine discrimination task they relied on responses from neurons tuned to flanking orientations. By dissociating the effects of perceptual load from other factors relating to differences in the physical stimuli or specific task requirements (e.g. fine vs. coarse discriminations), for example by using a dual-task paradigm, it could be determined whether the adaptive allocation of attentional gain relies on the availability of attentional resources and would therefore be affected by manipulations of perceptual load.

Likewise, it remains an open question how the reported changes in tuning due to limited perceptual resources develop over time. By combining two of the approaches used in this thesis (orientation noise masking and pattern backward masking) it might be possible to estimate the exact time course of changes in tuning by obtaining psychophysical orientation tuning functions under low and high load at various SOAs. First of all, if any change in tuning is observed depending on the SOA between target and mask, this would provide further evidence for the involvement of recurrent feedback in shaping the population orientation tuning profile. Secondly, by measuring changes in the temporal pattern of orientation selectivity with perceptual load the results could elucidate the cause of the extended suppression by the mask under high compared to low load observed in the current thesis (e.g. if tuning would be found to remain broad at later SOAs under high load, while it sharpens again when performance improves at earlier SOAs under low load).

6.6 Final conclusion

Inattention can have profound effects on visual perception that often go beyond the subjective expectations of observers. This thesis aimed to contribute to our understanding of the underlying mechanisms of reduced perception in inattention, specifically by investigating the role of the level of perceptual load imposed by the attended task. The results extend the framework of perceptual load theory to incorporate the modulation of contextual interactions among unattended stimuli and demonstrate how the level of load not only determines the magnitude of responses corresponding to unattended stimuli, but also changes the quality or fidelity of such representations by influencing selectivity and efficient separation of signal and noise in the spatial and temporal domains.

References

Adrian, E. D. (1926). The impulses produced by sensory nerve endings Part I. *The Journal of Physiology*, *61*(1), 49-72.

Albrecht, D. G., & Hamilton, D. B. (1982). Striate cortex of monkey and cat: contrast response function. *Journal of Neurophysiology*.

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

Bachmann, T. (1984). The process of perceptual retouch: Nonspecific afferent activation dynamics in explaining visual masking. *Perception & Psychophysics*, *35*(1), 69-84.

Bachmann, T., & Allik, J. (1976). Integration and interruption in the masking of form by form. *Perception*, *5*(1), 79-97.

Bahrami, B., Carmel, D., Walsh, V., Rees, G., & Lavie, N. (2008). Unconscious orientation processing depends on perceptual load. *Journal of Vision*, *8*(3).

Bahrami, B., Lavie, N., & Rees, G. (2007). Attentional load modulates responses of human primary visual cortex to invisible stimuli. *Current Biology*, *17*(6), 509-513.

Baldassi, S., & Burr, D. C. (2000). Feature-based integration of orientation signals in visual search. *Vision Research*, *40*(10), 1293-1300.

Baldassi, S., & Verghese, P. (2005). Attention to locations and features: Different top-down modulation of detector weights. *Journal of Vision*, *5*(6).

Bashinski, H. S., & Bacharach, V. R. (1980). Enhancement of perceptual sensitivity as the result of selectively attending to spatial locations. *Perception & Psychophysics*, *28*(3), 241-248.

Beck, D. M., & Lavie, N. (2005). Look here but ignore what you see: effects of distractors at fixation. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(3), 592.

Bergen, J. R., & Julesz, B. (1983). Parallel versus serial processing in rapid pattern discrimination. *Nature, 303*(5919), 696-698.

Blake, R., & Holopigian, K. (1985). Orientation selectivity in cats and humans assessed by masking. *Vision Research*, *25*(10), 1459-1467.

Blake, R., Holopigian, K., & Jauch, M. (1985). Another visual illusion involving orientation. *Vision Research*, *25*(10), 1469-1476.

Blakemore, C., & Tobin, E. A. (1972). Lateral inhibition between orientation detectors in the cat's visual cortex. *Experimental Brain Research*, *15*(4), 439-440.

Boehler, C. N., Tsotsos, J. K., Schoenfeld, M. A., Heinze, H. J., & Hopf, J. M. (2009). The center-surround profile of the focus of attention arises from recurrent processing in visual cortex. *Cerebral Cortex*, *19*(4), 982-991.

Bowen, R. W., & Wilson, H. R. (1994). A two-process analysis of pattern masking. *Vision Research*, *34*(5), 645-657.

Boyer, J., & Ro, T. (2007). Attention attenuates metacontrast masking. *Cognition*, *104*(1), 135-149.

Boynton, G. M., Demb, J. B., Glover, G. H., & Heeger, D. J. (1999). Neuronal basis of contrast discrimination. *Vision Research*, *39*(2), 257-269.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial vision*, *10*(4), 433-436.

Breitmeyer, B. G. (1984). *Visual masking: An integrative approach*. Oxford, UK: Oxford University Press.

Breitmeyer, B. G., & Ganz, L. (1976). Implications of sustained and transient channels for theories of visual pattern masking, saccadic suppression, and information processing. *Psychological Review*, *83*(1), 1.

Breitmeyer, B. G., & Ogmen, H. (2000). Recent models and findings in visual backward masking: A comparison, review, and update. *Perception & Psychophysics*, *62*(8), 1572-1595.

Broadbent, D. E. (1958). *Perception and communication*. New York: Oxford University Press.

Buracas, G. T., & Albright, T. D. (2009). Modulation of neuronal responses during covert search for visual feature conjunctions. *Proceedings of the National Academy of Sciences*, *106*(39), 16853-16858.

Busse, L., Katzner, S., Tillmann, C., & Treue, S. (2008). Effects of attention on perceptual direction tuning curves in the human visual system. *Journal of Vision*, *8*(9).

Cameron, E. L., Tai, J. C., & Carrasco, M. (2002). Covert attention affects the psychometric function of contrast sensitivity. *Vision Research*, *42*(8), 949-967.

Cameron, E. L., Tai, J. C., Eckstein, M. P., & Carrasco, M. (2004). Signal detection theory applied to three visual search tasks - identification, yes/no detection and localization. *Spatial Vision*, *17*(4), 295-326.

Caputo, G. (1998). Texture brightness filling-in. *Vision Research*, *38*(6), 841-851.

Carmel, D., Saker, P., Rees, G., & Lavie, N. (2007). Perceptual load modulates conscious flicker perception. *Journal of Vision*, *7*(14).

Carmel, D., Thorne, J. D., Rees, G., & Lavie, N. (2011). Perceptual load alters visual excitability. *Journal of Experimental Psychology: Human Perception and Performance*, *37*(5), 1350.

Carpenter, R. H. S., & Blakemore, C. (1973). Interactions between orientations in human vision. *Experimental Brain Research*, *18*(3), 287-303.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484-1525.

Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature Neuroscience*, *7*(3), 308-313.

Carrasco, M., Penpeci-Talgar, C., & Eckstein, M. (2000). Spatial covert attention increases contrast sensitivity across the CSF: Support for signal enhancement. *Vision Research*, *40*(10), 1203-1215.

Carrasco, M., Williams, P. E., & Yeshurun, Y. (2002). Covert attention increases spatial resolution with or without masks: Support for signal enhancement. *Journal of Vision*, *2*(6).

Cartwright-Finch, U., & Lavie, N. (2007). The role of perceptual load in inattentional blindness. *Cognition*, *102*(3), 321-340.

Cheesman, J., & Merikle, P. M. (1986). Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, *40*(4), 343.

Chelazzi, L., Duncan, J., Miller, E. K., & Desimone, R. (1998). Responses of neurons in inferior temporal cortex during memory-guided visual search. *Journal of Neurophysiology*, *80*(6), 2918-2940.

Clifford, C. W., & Harris, J. A. (2005). Contextual modulation outside of awareness. *Current Biology*, *15*(6), 574-578.

Cohen, M. R., & Maunsell, J. H. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience*, *12*(12), 1594-1600.

Crist, R. E., Li, W., & Gilbert, C. D. (2001). Learning to see: experience and attention in primary visual cortex. *Nature Neuroscience*, *4*(5), 519-525.

Debner, J. A., & Jacoby, L. L. (1994). Unconscious perception: attention, awareness, and control. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(2), 304.

Deneve, S., Latham, P. E., & Pouget, A. (1999). Reading population codes: a neural implementation of ideal observers. *Nature Neuroscience*, *2*(8), 740-745.

De Valois, R. L., & De Valois, K. K., (1988). *Spatial Vision*. Oxford, UK: Oxford University Press.

Di Lollo, V., Bischof, W. F., & Dixon, P. (1993). Stimulus-onset asynchrony is not necessary for motion perception or metacontrast masking. *Psychological Science*, *4*(4), 260-263.

Di Lollo, V., Enns, J. T., & Rensink, R. A. (2000). Competition for consciousness among visual events: the psychophysics of reentrant visual processes. *Journal of Experimental Psychology: General*, *129*(4), 481.

Dosher, B. A., & Lu, Z. L. (2000a). Noise exclusion in spatial attention.*Psychological Science*, *11*(2), 139-146.

Dosher, B. A., & Lu, Z. L. (2000b). Mechanisms of perceptual attention in precuing of location. *Vision Research*, *40*(10), 1269-1292.

Dresp, B. (1993). Bright lines and edges facilitate the detection of small light targets. *Spatial Vision*, *7*(3), 213-225.

Driver, J. (2001). A selective review of selective attention research from the past century. *British Journal of Psychology*, *92*(1), 53-78.

Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review*, *87*(3), 272.

Eckstein, M. P., Peterson, M. F., Pham, B. T., & Droll, J. A. (2009). Statistical decision theory to relate neurons to behavior in the study of covert visual attention. *Vision Research*, *49*(10), 1097-1128.

Eckstein, M. P., Shimozaki, S. S., & Abbey, C. K. (2002). The footprints of visual attention in the Posner cueing paradigm revealed by classification images. *Journal of Vision*, *2*(1), 25-45.

Eckstein, M. P., Thomas, J. P., Palmer, J., & Shimozaki, S. S. (2000). A signal detection model predicts the effects of set size on visual search accuracy for feature, conjunction, triple conjunction, and disjunction displays. *Perception & Psychophysics*, *62*(3), 425-451.

Enns, J. T. (2004). Object substitution and its relation to other forms of visual masking. *Vision Research*, *44*(12), 1321-1331.

Enns, J. T., & Di Lollo, V. (1997). Object substitution: A new form of masking in unattended visual locations. *Psychological Science*, *8*(2), 135-139.

Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, *392*(6676), 598-601.

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, *16*(1), 143-149.

Felsen, G., Shen, Y. S., Yao, H., Spor, G., Li, C., & Dan, Y. (2002). Dynamic modification of cortical orientation tuning mediated by recurrent connections. *Neuron*, *36*(5), 945-954.

Felsen, G., Touryan, J., & Dan, Y. (2005). Contextual modulation of orientation tuning contributes to efficient processing of natural stimuli. *Network: Computation in Neural Systems*, *16*(2-3), 139-149.

Ferster, D., Chung, S., & Wheat, H. (1996). Orientation selectivity of thalamic input to simple cells of cat visual cortex. *Nature*, *380*(6571), 249-252.

Ferster, D., & Miller, K. D. (2000). Neural mechanisms of orientation selectivity in the visual cortex. *Annual Review of Neuroscience*, *23*(1), 441-471.

Foley, J. M., & Schwarz, W. (1998). Spatial attention: Effect of position uncertainty and number of distractor patterns on the threshold-versus-contrast function for contrast discrimination. *Journal of the Optical Society of America A*, *15*(5), 1036-1047.

Forster, S., & Lavie, N. (2007). High Perceptual Load Makes Everybody Equal Eliminating Individual Differences in Distractibility With Load. *Psychological Science*, *18*(5), 377-381.

Forster, S., & Lavie, N. (2008). Failures to ignore entirely irrelevant distractors: the role of load. *Journal of Experimental Psychology: Applied*, *14*(1), 73.

Freeman, E., Driver, J., Sagi, D., & Zhaoping, L. (2003). Top-down modulation of lateral interactions in early vision: Does attention affect integration of the whole or just perception of the parts? *Current Biology*, *13*(11), 985-989.

Freeman, E., Sagi, D., & Driver, J. (2001). Lateral interactions between targets and flankers in low-level vision depend on attention to the flankers. *Nature Neuroscience*, *4*(10), 1032-1036.

Freeman, E. D., Sagi, D., & Driver, J. (2004). Configuration-specific attentional modulation of flanker target lateral interactions. *Perception*, *33*(2), 181-194.

Gandhi, S. P., Heeger, D. J., & Boynton, G. M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proceedings of the National Academy of Sciences*, *96*(6), 3314-3319.

Gilbert, C. D., & Li, W. (2013). Top-down influences on visual processing. *Nature Reviews Neuroscience*, *14*(5), 350-363.

Gilbert, C. D., & Wiesel, T. N. (1989). Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *The Journal of Neuroscience*, *9*(7), 2432-2442.

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.

Graham, N. V. S. (2001). *Visual pattern analyzers (Oxford Psychology Series: No. 16).* Oxford, UK: Oxford University Press.

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics (Vol. 1).* New York, NY: Wiley.

Grill-Spector, K., & Malach, R. (2001). fMR-adaptation: a tool for studying the functional properties of human cortical neurons. *Acta Psychologica*, *107*(1), 293-321.

Gross, C. G., Rocha-Miranda, C. E., & Bender, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the Macaque. *Journal of Neurophysiology*.

Growney, R., Weisstein, N., & Cox, S. I. (1977). Metacontrast as a function of spatial separation with narrow line targets and masks. *Vision Research*, *17*(10), 1205-1210.

Hopf, J. M., Boehler, C. N., Luck, S. J., Tsotsos, J. K., Heinze, H. J., & Schoenfeld, M. A. (2006). Direct neurophysiological evidence for spatial suppression surrounding the focus of attention in vision. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(4), 1053-1058.

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.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, *160*(1), 106.

Ito, M., & Gilbert, C. D. (1999). Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron*, *22*(3), 593-604.

Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience, 9*(5), 690-696.

Jonides, J., & Yantis, S. (1988). Uniqueness of abrupt visual onset in capturing attention. *Perception & Psychophysics*, *43*(4), 346-354.

Joo, S. J., Boynton, G. M., & Murray, S. O. (2012). Long-range, pattern-dependent contextual effects in early human visual cortex. *Current Biology*, *22*(9), 781-786.

Kahneman, D. (1968). Method, findings, and theory in studies of visual masking. *Psychological Bulletin*, *70*(6p1), 404.

Kamitani, Y., & Tong, F. (2006). Decoding seen and attended motion directions from activity in the human visual cortex. *Current Biology*, *16*(11), 1096-1102.

Kapadia, M. K., Ito, M., Gilbert, C. D., & Westheimer, G. (1995). Improvement in visual sensitivity by changes in local context: parallel studies in human observers and in V1 of alert monkeys. *Neuron*, *15*(4), 843-856.

Kapadia, M. K., Westheimer, G., & Gilbert, C. D. (1999). Dynamics of spatial summation in primary visual cortex of alert monkeys. *Proceedings of the National Academy of Sciences*, *96*(21), 12073-12078.

Kapadia, M. K., Westheimer, G., & Gilbert, C. D. (2000). Spatial distribution of contextual interactions in primary visual cortex and in visual perception. *Journal of Neurophysiology*, *84*(4), 2048-2062.

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.

Keysers, C., & Perrett, D. I. (2002). Visual masking and RSVP reveal neural competition. *Trends in Cognitive Sciences*, *6*(3), 120-125.

Kim, Y. J., Grabowecky, M., Paller, K. A., Muthu, K., & Suzuki, S. (2006). Attention induces synchronization-based response gain in steady-state visual evoked potentials. *Nature Neuroscience*, *10*(1), 117-125.

Kinchla, R. A. (1992). Attention. *Annual Review of Psychology,* 43, 711-43.

Kinoshita, M., Gilbert, C. D., & Das, A. (2009). Optical imaging of contextual interactions in V1 of the behaving monkey. *Journal of Neurophysiology*, *102*(3), 1930-1944.

Kinsbourne, M., & Warrington, E. K. (1962). The effect of an after-coming random pattern on the perception of brief visual stimuli. *Quarterly Journal of Experimental Psychology*, *14*(4), 223-234.

Kovacs, G., Vogels, R., & Orban, G. A. (1995). Cortical correlate of pattern backward masking. *Proceedings of the National Academy of Sciences*, *92*(12), 5587-5591.

Lamme, V. A. (2000). Neural mechanisms of visual awareness: A linking proposition. *Brain and Mind*, *1*(3), 385-406.

Lamme, V. A., Rodriguez-Rodriguez, V., & Spekreijse, H. (1999). Separate processing dynamics for texture elements, boundaries and surfaces in primary visual cortex of the macaque monkey. *Cerebral Cortex*, *9*(4), 406-413.

Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, *23*(11), 571-579.

Lamme, V. A., Super, H., & Spekreijse, H. (1998). Feedforward, horizontal, and feedback processing in the visual cortex. *Current Opinion in Neurobiology*, *8*(4), 529-535.

Lamme, V. A., Zipser, K., & Spekreijse, H. (2002). Masking interrupts figure-ground signals in V1. *Journal of Cognitive Neuroscience*, *14*(7), 1044-1053.

Lavie, N. (1995). Perceptual load as a necessary condition for selective attention. *Journal of Experimental Psychology: Human Perception and Performance*, *21*(3), 451.

Lavie, N. (2000). Selective attention and cognitive control: Dissociating attentional functions through different types of load. *Attention and Performance XVIII*, 175-194.

Lavie, N. (2005). Distracted and confused?: Selective attention under load. *Trends in Cognitive Sciences*, *9*(2), 75-82.

Lavie, N. (2010). Attention, distraction, and cognitive control under load. *Current Directions in Psychological Science*, *19*(3), 143-148.

Lavie, N., & Cox, S. (1997). On the efficiency of visual selective attention: Efficient visual search leads to inefficient distractor rejection. *Psychological Science*, *8*(5), 395-396.

Lavie, N., & De Fockert, J. W. (2003). Contrasting effects of sensory limits and capacity limits in visual selective attention. *Perception & Psychophysics*, *65*(2), 202-212.

Lavie, N., & Fox, E. (2000). The role of perceptual load in negative priming. *Journal of Experimental Psychology: Human Perception and Performance*, *26*(3), 1038.

Lavie, N., Hirst, A., de Fockert, J. W., & Viding, E. (2004). Load theory of selective attention and cognitive control. *Journal of Experimental Psychology: General*, *133*(3), 339.

Lee, D. K., Itti, L., Koch, C., & Braun, J. (1999). Attention activates winner-take-all competition among visual filters. *Nature Neuroscience*, *2*(4), 375-381.

Legge, G. E., & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America*, *70*(12), 1458-1471.

Lennie, P. (2003). The cost of cortical computation. *Current Biology*, *13*(6), 493-497.

Li, W., Thier, P., & Wehrhahn, C. (2000). Contextual influence on orientation discrimination of humans and responses of neurons in V1 of alert monkeys. *Journal of Neurophysiology*, *83*(2), 941-954.

Ling, S., & Blake, R. (2009). Suppression during binocular rivalry broadens orientation tuning. *Psychological Science*, *20*(11), 1348-1355.

Ling, S., & Carrasco, M. (2006a). Sustained and transient covert attention enhance the signal via different contrast response functions. *Vision Research*, *46*(8), 1210-1220.

Liu, T., Larsson, J., & Carrasco, M. (2007). Feature-based attention modulates orientation-selective responses in human visual cortex. *Neuron*, *55*(2), 313-323.

Ling, S., Liu, T., & Carrasco, M. (2009). How spatial and feature-based attention affect the gain and tuning of population responses. *Vision Research*, *49*(10), 1194-1204.

Ling, S., Pearson, J., & Blake, R. (2009). Dissociation of neural mechanisms underlying orientation processing in humans. *Current Biology*, *19*(17), 1458-1462.

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

Lu, Z. L., & Dosher, B. A. (2000). Spatial attention: Different mechanisms for central and peripheral temporal precues?. *Journal of Experimental Psychology: Human Perception and Performance*, *26*(5), 1534.

Lu, Z. L., & Dosher, B. A. (2004). Spatial attention excludes external noise without changing the spatial frequency tuning of the perceptual template. *Journal of Vision*, *4*(10).

Lu, Z. L., Lesmes, L. A., & Dosher, B. A. (2002). Spatial attention excludes external noise at the target location. *Journal of Vision*, *2*(4).

Lu, Z. L., Liu, C. Q., & Dosher, B. A. (2000). Attention mechanisms for multi-location first-and second-order motion perception. *Vision Research*, *40*(2), 173-186.

Luckiesh, M. (1922). *Visual illusions: Their causes, characteristics and applications (Vol. 2)*. Dover, NY: D. Van Nostrand Company.

Macdonald, J. S., & Lavie, N. (2008). Load induced blindness. *Journal of Experimental Psychology: Human Perception and Performance*, *34*(5), 1078.

Mack, A., & Rock, I. (1998). *Inattentional blindness*. Cambridge, MA: The MIT Press.

Macknik, S. L., & Livingstone, M. S. (1998). Neuronal correlates of visibility and invisibility in the primate visual system. *Nature Neuroscience*, *1*(2), 144-149.

Majaj, N. J., Pelli, D. G., Kurshan, P., & Palomares, M. (2002). The role of spatial frequency channels in letter identification. *Vision Research, 42*(9), 1165-1184.

Mareschal, I., Solomon, J., & Morgan, M. (2009). Endogenous attention can reduce the tilt illusion, but not crowding. *Journal of Vision*, *9*(8), 126-126.

Mariño, J., Schummers, J., Lyon, D. C., Schwabe, L., Beck, O., Wiesing, P., Obermayer, K. & Sur, M. (2005). Invariant computations in local cortical networks with balanced excitation and inhibition. *Nature Neuroscience*, *8*(2), 194-201.

Martinez-Trujillo, J. C., & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology*, *14*(9), 744-751.

Martinez-Trujillo, J. C., Treue, S. (2005). Attentional modulation of apparent stimulus contrast. In L. Itti, G. Rees, & J. K. Tsotsos (Eds.), *Neurobiology of attention* (p. 428). London, UK: Elsevier Academic Press.

Maunsell, J. H., & Cook, E. P. (2002). The role of attention in visual processing. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, *357*(1424), 1063-1072.

Maunsell, J. H., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, *29*(6), 317-322.

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

Michaels, C. F., & Turvey, M. T. (1979). Central sources of visual masking: Indexing structures supporting seeing at a single, brief glance. *Psychological Research*, *41*(1), 1-61.

Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, *229*(4715), 782-784.

Morgan, M. J., Ward, R. M., & Castet, E. (1998). Visual search for a tilted target: Tests of spatial uncertainty models. *The Quarterly Journal of Experimental Psychology: Section A*, *51*(2), 347-370.

Morrone, M. C., Denti, V., & Spinelli, D. (2004). Different attentional resources modulate the gain mechanisms for color and luminance contrast. *Vision Research*, *44*(12), 1389-1401.

Motter, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *Journal of Neurophysiology*, *70*(3), 909-919.

Muggleton, N., Lamb, R., Walsh, V., & Lavie, N. (2008). Perceptual load modulates visual cortex excitability to magnetic stimulation. *Journal of Neurophysiology*, *100*(1), 516-519.

Muise, J. G., LeBlanc, R. S., Lavoie, M. E., & Arsenault, A. S. (1991). Two-stage model of visual backward masking: Sensory transmission and accrual of effective information as a function of target intensity and similarity. *Perception & Psychophysics*, *50*(3), 197-204.

Müller, J. R., Metha, A. B., Krauskopf, J., & Lennie, P. (2003). Local signals from beyond the receptive fields of striate cortical neurons. *Journal of Neurophysiology*, *90*(2), 822-831.

Murray, R. F., Sekuler, A. B., & Bennett, P. J. (2003). A linear cue combination framework for understanding selective attention. *Journal of Vision*, *3*(2).

Naccache, L., Blandin, E., & Dehaene, S. (2002). Unconscious masked priming depends on temporal attention. *Psychological Science*, *13*(5), 416-424.

Nachmias, J. (1967). Effect of exposure duration on visual contrast sensitivity with square-wave gratings. *Journal of the Optical Society of America*, *57*(3), 421-427.

Nachmias, J. (2002). Contrast discrimination with and without spatial uncertainty. *Vision Research*, *42*(1), 41-48.

Navalpakkam, V., & Itti, L. (2007). Search goal tunes visual features optimally. *Neuron*, *53*(4), 605-617.

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

Neri, P. (2004). Attentional effects on sensory tuning for single-feature detection and double-feature conjunction. *Vision Research*, *44*(26), 3053-3064.

Neri, P., & Levi, D. M. (2006). Receptive versus perceptive fields from the reverse-correlation viewpoint. *Vision Research*, *46*(16), 2465-2474.

Norman, D. A., & Bobrow, D. G. (1975). On data-limited and resource-limited processes. *Cognitive Psychology*, *7*(1), 44-64.

O'Connor, D. H., Fukui, M. M., Pinsk, M. A., & Kastner, S. (2002). Attention modulates responses in the human lateral geniculate nucleus. *Nature Neuroscience*, *5*(11), 1203-1209.

Palmer, J. (1994). Set-size effects in visual search: The effect of attention is independent of the stimulus for simple tasks. *Vision Research*, *34*(13), 1703-1721.

Palmer, J., Verghese, P., & Pavel, M. (2000). The psychophysics of visual search. *Vision Research*, *40*(10), 1227-1268.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437-442.

Pestilli, F., & Carrasco, M. (2005). Attention enhances contrast sensitivity at cued and impairs it at uncued locations. *Vision Research*, *45*(14), 1867-1875.

Pinsk, M. A., Doniger, G. M., & Kastner, S. (2004). Push-pull mechanism of selective attention in human extrastriate cortex. *Journal of Neurophysiology*, *92*(1), 622-629.

Polat, U., Mizobe, K., Pettet, M. W., Kasamatsu, T., & Norcia, A. M. (1998). Collinear stimuli regulate visual responses depending on cell's contrast threshold. *Nature*, *391*(6667), 580-584.

Polat, U., & Sagi, D. (1993). Lateral interactions between spatial channels: suppression and facilitation revealed by lateral masking experiments. *Vision Research*, *33*(7), 993-999.

Polat, U., & Sagi, D. (1994). The architecture of perceptual spatial interactions. *Vision Research*, *34*(1), 73-78.

Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience*, *1*(2), 125-132.

Pouget, A., Dayan, P., & Zemel, R. S. (2003). Inference and computation with population codes. *Annual Review of Neuroscience*, *26*(1), 381-410.

Priebe, N. J., & Ferster, D. (2008). Inhibition, spike threshold, and stimulus selectivity in primary visual cortex. *Neuron*, *57*(4), 482-497.

Ramachandran, V. S., & Cobb, S. (1995). Visual attention modulates metacontrast masking. *Nature*, *373*(6509), 66-68.

Rees, G., Frith, C. D., & Lavie, N. (1997). Modulating irrelevant motion perception by varying attentional load in an unrelated task. *Science*, *278*(5343), 1616-1619.

Regan, D., & Beverley, K. I. (1985). Postadaptation orientation discrimination. *Journal of the Optical Society of America A*, *2*(2), 147-155.

Reynolds, J.H. (2005). Visual cortical circuits and spatial attention. In L. Itti, G. Rees, & J. K. Tsotsos (Eds.), *Neurobiology of attention* (p. 42-49). London, UK: Elsevier Academic Press.

Reynolds, J. H., & Chelazzi, L. (2004). Attentional modulation of visual processing. *Annual Review of Neuroscience*, *27*, 611-647.

Reynolds, J. H., Chelazzi, L., & Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *The Journal of Neuroscience*, *19*(5), 1736-1753.

Reynolds, J. H., & Desimone, R. (2003). Interacting roles of attention and visual salience in V4. *Neuron*, *37*(5), 853-863.

Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron*, *61*(2), 168-185.

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

Ringach, D. L., Hawken, M. J., & Shapley, R. (1997). Dynamics of orientation tuning in macaque primary visual cortex. *Nature*, *387*(6630), 281-284.

Rolls, E. T., Tovée, M. J., & Panzeri, S. (1999). The neurophysiology of backward visual masking: information analysis. *Journal of Cognitive Neuroscience*, *11*(3), 300-311.

Roper, Z. J., Cosman, J. D., & Vecera, S. P. (2013). Perceptual Load Corresponds With Factors Known to Influence Visual Search. *Journal of Experimental Psychology: Human Perception and Performance*, [E-pub ahead of print]. doi: 10.1037/a0031616.

Sagi, D., & Julesz, B. (1985). Detection versus discrimination of visual orientation. *Perception*, *14*, 619-628.

Scharf, B., & Lefton, L. A. (1970). Backward and forward masking as a function of stimulus and task parameters. *Journal of Experimental Psychology*, *84*(2), 331.

Scheerer, E. (1973). Integration, interruption and processing rate in visual backward masking. *Psychologische Forschung*, *36*(1), 71-93.

Schiller, P. H. (1966). Forward and backward masking as a function of relative overlap and intensity of test and masking stimuli. *Perception & Psychophysics*, *1*(3), 161-164.

Schwartz, O., & Coen-Cagli, R. (2013). Visual attention and flexible normalization pools. *Journal of Vision*, *13*(1).

Schwartz, S., Vuilleumier, P., Hutton, C., Maravita, A., Dolan, R. J., & Driver, J. (2005). Attentional load and sensory competition in human vision: modulation of fMRI responses by load at fixation during task-irrelevant stimulation in the peripheral visual field. *Cerebral Cortex, 15*(6), 770-786.

Sclar, G., & Freeman, R. D. (1982). Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Experimental Brain Research*, *46*(3), 457-461.

Sclar, G., Lennie, P., & DePriest, D. D. (1989). Contrast adaptation in striate cortex of macaque. *Vision Research*, *29*(7), 747-755.

Sclar, G., Maunsell, J. H., & Lennie, P. (1990). Coding of image contrast in central visual pathways of the macaque monkey. *Vision Research*, *30*(1), 1-10.

Scolari, M., & Serences, J. T. (2009). Adaptive allocation of attentional gain. *The Journal of Neuroscience*, *29*(38), 11933-11942.

Sengpiel, F., Sen, A., & Blakemore, C. (1997). Characteristics of surround inhibition in cat area 17. *Experimental Brain Research*, *116*(2), 216-228.

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.

Shadlen, M. N., Britten, K. H., Newsome, W. T., & Movshon, J. A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *The Journal of Neuroscience, 16*(4), 1486-1510.

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *The Journal of neuroscience*, *18*(10), 3870-3896.

Shapley, R., Hawken, M., & Ringach, D. L. (2003). Dynamics of orientation selectivity in the primary visual cortex and the importance of cortical inhibition. *Neuron, 38*(5), 689-699

Shaw, M.L. (1984). Division of attention among spatial locations: A fundamental difference between detection of letters and detection of luminance. In H. Bouma & D.G. Bouwhuis (Eds.), *Attention and Performance X* (p. 109-120). Hillsdale, NJ: Erlbaum.

Sheinberg, D. L., & Logothetis, N. K. (2001). Noticing familiar objects in real world scenes: the role of temporal cortical neurons in natural vision. *The Journal of Neuroscience*, *21*(4), 1340-1350.

Shelley-Tremblay, J., & Mack, A. (1999). Metacontrast masking and attention. *Psychological Science*, *10*(6), 508-515.

Shiu, L. P., & Pashler, H. (1994). Negligible effect of spatial precuing on identification of single digits. *Journal of Experimental Psychology: Human Perception and Performance*, *20*(5), 1037.

Shou, T. D., & Leventhal, A. G. (1989). Organized arrangement of orientation-sensitive relay cells in the cat's dorsal lateral geniculate nucleus. *The Journal of Neuroscience*, *9*(12), 4287-4302.

Solomon, J. A., Lavie, N., & Morgan, M. J. (1997). Contrast discrimination function: Spatial cuing effects. *Journal of the Optical Society of America A*, *14*(9), 2443-2448.

Solomon, J. A., & Pelli, D. G. (1994). The visual filter mediating letter identification. *Nature*, *369*(6479), 395-397.

Somers, D. C., Nelson, S. B., & Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *The Journal of Neuroscience*, *15*(8), 5448-5465.

Sompolinsky, H., & Shapley, R. (1997). New perspectives on the mechanisms for orientation selectivity. *Current Opinion in Neurobiology*, *7*(4), 514-522.

Spencer, T. J. (1969). Some effects of different masking stimuli on iconic storage. *Journal of Experimental Psychology*, *81*(1), 132.

Spencer, T. J., & Shuntich, R. (1970). Evidence for an interruption theory of backward masking. *Journal of Experimental Psychology*, *85*(2), 198.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, *74*(11), 1.

Sperling, G. (1963). A model for visual memory tasks. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *5*(1), 19-31.

Sperling, G. (1967). Successive approximations to a model for short term memory. *Acta Psychologica*, *27*, 285-292.

Sperling G., & Dosher, B. A. (1986). Strategy and optimization in human information processing. In K. Boff, L. Kaufman & J. Thomas (Eds.), *Handbook of perception and performance Vol. 1* (p. 1-85). New York, NY: Wiley.

Spitzer, H., Desimone, R., & Moran, J. (1988). Increased attention enhances both behavioral and neuronal performance. *Science*, *240*(4850), 338-340.

Spivey, M. J., & Spirn, M. J. (2000). Selective visual attention modulates the direct tilt aftereffect. *Perception & Psychophysics*, *62*(8), 1525-1533.

Sundberg, K. A., Mitchell, J. F., & Reynolds, J. H. (2009). Spatial attention modulates center-surround interactions in macaque visual area v4. *Neuron*, *61*(6), 952-963.

Talgar, C. P., Pelli, D. G., & Carrasco, M. (2004). Covert attention enhances letter identification without affecting channel tuning. *Journal of Vision*, *4*(1).

Tata, M. S. (2002). Attend to it now or lose it forever: Selective attention, metacontrast masking, and object substitution. *Perception & Psychophysics, 64*(7), 1028-1038.

Tata, M. S., & Giaschi, D. E. (2004). Warning: Attending to a mask may be hazardous to your perception. *Psychonomic Bulletin & Review*, *11*(2), 262-268.

Thompson, K. G., & Schall, J. D. (1999). The detection of visual signals by macaque frontal eye field during masking. *Nature Neuroscience, 2*(3), 283-288.

Torralbo, A., & Beck, D. M. (2008). Perceptual-load-induced selection as a result of local competitive interactions in visual cortex. *Psychological Science, 19*(10), 1045-1050.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97-136.

Treue, S. (2001). Neural correlates of attention in primate visual cortex. *Trends in Neurosciences*, *24*(5), 295-300.

Treue, S., & Martinez-Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature, 399*(6736), 575-579.

Troyer, T. W., Krukowski, A. E., Priebe, N. J., & Miller, K. D. (1998). Contrast-invariant orientation tuning in cat visual cortex: thalamocortical input tuning and correlation-based intracortical connectivity. *The Journal of Neuroscience, 18*(15), 5908-5927.

Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience, 8*(8), 1096-1101.

Tsushima, Y., Sasaki, Y., & Watanabe, T. (2006). Greater disruption due to failure of inhibitory control on an ambiguous distractor. *Science*, *314*(5806), 1786-1788.

Tsushima, Y., Seitz, A. R., & Watanabe, T. (2008). Task-irrelevant learning occurs only when the irrelevant feature is weak. *Current Biology*, *18*(12), 516-517.

Turvey, M. T. (1973). On peripheral and central processes in vision: Inferences from an information-processing analysis of masking with patterned stimuli. *Psychological Review*, *80*, 1-52.

Tzvetanov, T., Womelsdorf, T., Niebergall, R., & Treue, S. (2006). Feature-based attention influences contextual interactions during motion repulsion. *Vision Research*, *46*(21), 3651-3658.

Uttal, W. R. (1970). On the physiological basis of masking with dotted visual noise. *Perception & Psychophysics*, *7*(6), 321-327.

Vanrullen, R., & Thorpe, S. J. (2001). The time course of visual processing: from early perception to decision-making. *Journal of Cognitive Neuroscience*, *13*(4), 454-461.

Watson, A. B., & Pelli, D. G. (1983). QUEST: A Bayesian adaptive psychometric method. *Perception & Psychophysics*, *33*(2), 113-120.

Wehrhahn, C., Li, W., & Westheimer, G. (1996). Patterns that impair discrimination of line orientation in human vision. *Perception*, *25*, 1053-1064.

Weisstein, N. (1971). W-shaped and U-shaped functions obtained for monoptic and dichoptic disk-disk masking. *Perception & Psychophysics*, *9*(3), 275-278.

Williford, T., & Maunsell, J. H. (2006). Effects of spatial attention on contrast response functions in macaque area V4. *Journal of Neurophysiology*, *96*(1), 40-54.

Wolfe J. M., (1999). Inattentional amnesia. In V. Coltheart (Ed.), *Fleeting Memories: Cognition of Brief Visual Stimuli* (p. 71-94). Cambridge, MA: MIT Press.

Yantis, S. (1993). Stimulus-driven attentional capture. *Current Directions in Psychological Science, 2*(5), 156-161.

Yeshurun, Y., & Carrasco, M. (1998). Attention improves or impairs visual performance by enhancing spatial resolution. *Nature, 396*(6706), 72-75.

Yi, D. J., Woodman, G. F., Widders, D., Marois, R., & Chun, M. M. (2004). Neural fate of ignored stimuli: dissociable effects of perceptual and working memory load. *Nature Neuroscience, 7*(9), 992-996.

Zohary, E., Shadlen, M. N., & Newsome, W. T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, 370, 140-143.