SPATIALLY DIRECTED ATTENTION TO THE FEATURES OF OBJECTS: PRECISION AND CONTRIBUTIONS FROM SERIAL PROCESSING

CHRISTINA J. HOWARD

THESIS SUBMITTED FOR THE DEGREE OF PH.D
CARDIFF UNIVERSITY
2007

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

CARDIFF
UNIVERSITY
PRIFYSGOL
CAERDYDD

DECLARATION

This work has not previously been accepted in substance for any degree and is not concurrently submitted in candidature for any degree.

Signed ....C. Howard.............. (candidate) Date ....14/3/08.....

STATEMENT 1

This thesis is being submitted in partial fulfillment of the requirements for the degree of ....PhD...........(insert MCh, MD, MPhil, PhD etc, as appropriate)

Signed ....C. Howard.............. (candidate) Date ....14/3/08........

STATEMENT 2

This thesis is the result of my own independent work/investigation, except where otherwise stated. Other sources are acknowledged by explicit references.

Signed ....C. Howard.............. (candidate) Date ....14/3/08........

STATEMENT 3

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed ....C. Howard.............. (candidate) Date ....14/3/08........

SPATIALLY DIRECTED ATTENTION TO THE FEATURES OF OBJECTS: PRECISION AND CONTRIBUTIONS FROM SERIAL PROCESSING

SUMMARY OF THESIS

Two paradigms are presented to investigate the facilitatory effect of spatial cueing on perception of the features of objects.

In Chapter 2, performance was studied for continuous tracking of the positional and non-positional features of multiple objects. Observers attempted to report the last feature value of a queried object. The results demonstrated a progressive decline in the precision of the representation of tracked objects, for orientation, spatial period and position. The decline was apparent as the number of objects was increased from one to two to four, and occurred before previously suggested capacity limits based on a fixed number of objects.

Responses were more similar to past states of the queried object than to its last state, that is to say that responses exhibited perceptual lags. For orientation, and particularly for spatial period, perceptual lags increased as more objects were added to the attentional load. Analyses of error patterns broken down by likely confidence in responses suggested a contribution from serial processing. Further support for the notion that the lags were contributed to by serial processing, was suggested by two double report experiments presented in Chapter 3. Here, for spatial period tracking, processing appeared not to proceed purely in parallel. Why tracking performance may differ between features is discussed.

Four experiments in Chapter 4 rule out the possibilities that the differences in lags between orientation and spatial period were due to artefacts relating to the range of feature values.

Chapter 6 is an investigation into the benefits of attention in terms of noise reduction and sampling. Pre-cueing and dual task manipulations were used to compare the effects of attention with those of crowding, which has been demonstrated to increase internal noise. Spatially directed attention was found to reduce internal noise, consistent with the precision results of the previous chapters.

ACKNOWLEDGEMENTS

This thesis was undertaken at the School of Psychology, Cardiff University, under the supervision of Alex Holcombe, Petroc Sumner and Dylan Jones. Thank you to all three for your guidance and support throughout this project. Thank you particularly to Alex Holcombe for his continued interest, patience and feedback on the continuous monitoring work, to Petroc Sumner for his advice on the equivalent noise experiments, and to Dylan Jones for his support throughout my three years of study.

This work was supported by a studentship from the School of Psychology, Cardiff University funded by the EPSRC and the ESRC.

CHAPTER 1. GENERAL INTRODUCTION TO VISUAL ATTENTION............1

1.1 Models of general attention...1

1.2 Visual spatial attention...2

1.3 The necessity of attention for perception..2

1.4 Questions on attention in this thesis..4

CHAPTER 2. EXPERIMENTS 1 – 3: CONTINUOUS MONITORING OF THE FEATURES OF MULTIPLE OBJECTS..6

2.1 INTRODUCTION...6

2.1.1 Summary of literature covered in the Introduction.....................................6

2.1.2 Continuously monitoring positional features...6

2.1.3 The nature of the attentional limit..7

2.1.4 Further evidence against an attentional limit based on a fixed number of objects..8

2.1.5 Non-positional features..9

2.1.6 Non-positional features: visual short-term memory studies.......................10

2.1.7 Parallel and serial processing..12

2.2 METHOD FOR EXPERIMENTS 1 AND 2: ORIENTATION AND SPATIAL PERIOD TRACKING...14

2.2.1 Observers..16

2.2.2 Procedure..16

2.2.3 Design...19

2.2.4 Trajectories of Gabors through orientation or spatial period......................19

2.2.5 Trajectory parameters. Experiment 1: Orientation.....................................19

2.2.6 Trajectory parameters. Experiment 2: Spatial period....20

2.3 EXPERIMENT 3: POSITION TRACKING....21

2.3.1 Observers....22

2.3.2 Procedure....23

2.3.3 Design....25

2.3.4 Trajectories of discs....26

2.4 RESULTS AND DISCUSSION FOR EXPERIMENTS 1 – 3....*27*

2.4.1 Summary of set size effects on precision....32

2.4.2 Spatial bias in Experiment 3....33

2.4.3 Distracter effects and effect of target marker duration....34

2.4.4 Perceptual lag analysis....34

2.4.5 Serial processing: analysis by guessing likelihood....39

2.4.6 Results of analysis by guessing likelihood: contribution from serial processing....40

2.4.7 Eliminating explanations based on statistics of trajectories....47

2.4.8 Summary of guessing likelihood analysis: implications for serial processing....48

CHAPTER 3. SUMMARY OF GUESSING LIKELIHOOD ANALYSIS: IMPLICATIONS FOR SERIAL PROCESSING....**49**

3.1 INTRODUCTION....49

3.1.1 Overview of Experiments 4 and 5....49

3.1.2 Rationale for Experiments 4 and 5: distinguishing between serial and parallel processing....49

3.1.3 Worked example....50

3.2 METHOD FOR EXPERIMENTS 4 AND 5..52

3.2.1 Observers..52

3.2.2 Design...52

3.3 RESULTS AND DISCUSSION FOR EXPERIMENTS 4 AND 5.....................53

3.3.1 Summary of results from Experiments 4 & 5: Distinguishing between serial and parallel processing..56

CHAPTER 4. EXPERIMENTS 6 – 9: CONTROL EXPERIMENTS FOR POSSIBLE ARTEFACTS IN THE RANGE OF FEATURE VALUES.................................58

4.1 INTRODUCTION...58

4.1.1 Rationale for Experiments 6 – 9..58

4.1.2 Overview of Experiments 6 – 9..59

4.2 METHOD FOR EXPERIMENTS 6 – 9...60

4.2.1 Observers..61

4.3 RESULTS AND DISCUSSION OF EXPERIMENTS 6 – 9...........................62

4.3.1 Error analyses..62

4.3.2 Lag analyses..69

4.3.3 Differences in lag patterns between Experiment 1 and Experiments 6 – 9........74

4.3.4 Summary of findings in Experiments 6 – 9..75

CHAPTER 5. DISCUSSION OF CONTINUOUS MONITORING EXPERIMENTS IN CHAPTERS 2 – 4...76

5.1 SUMMARY OF RESULTS OF CHAPTERS 2 – 4..76

5.1.1 The nature of the attentional limit on continuously monitoring the features of multiple objects......76

5.1.2 Contribution from serial processing......77

5.2 ADVANTAGES OF THE CONTINUOUS MONITORING PARADIGM......78

5.3 ICONIC MEMORY AND VISUAL PERSISTENCE......79

5.4 RELATION TO PREVIOUS POSITIONAL FEATURE TRACKING STUDIES......80

5.5 RELATION TO PREVIOUS RESULTS RELATING ATTENTION TO PRECISION......82

5.6 PERCEPTUAL LAGS......84

5.6.1 Relation of perceptual lags to other phenomena in motion perception......86

5.6.2 Implications of perceptual lags......87

5.7 PARALLEL AND SERIAL PROCESSING......87

5.8 DIFFERENCES BETWEEN FEATURES......92

5.9 FUTURE DIRECTIONS......93

CHAPTER 6. EXPERIMENTS 10 AND 11: COMPARISON OF ATTENTION AND CROWDING: EFFECTS ON NOISE AND SAMPLING......95

6.1 INTRODUCTION......95

6.1.1 Feature integration and origins of attentional explanations of crowding......95

6.1.2 Support for an attentional explanation of crowding......97

6.1.3 Opponents to the attentional hypothesis......97

6.1.4 Attention, crowding and internal noise......99

6.1.5 Overview of experiments......102

6.2 EXPERIMENT 10 (DUAL TASK)..103

6.3 METHOD FOR EXPERIMENT 10..103

6.3.1 Observers..104

6.3.2 Procedure..104

6.3.3 The secondary task..105

6.3.4 The primary task..106

6.3.5 Response arrangements..106

6.3.6 Design..107

6.4 RESULTS FOR EXPERIMENT 10 (DUAL TASK)..107

6.5 EXPERIMENT 11 (PRE-CUEING)..111

6.6 METHOD FOR EXPERIMENT 11..111

6.6.1 Observers..111

6.6.2 Procedure..112

6.6.3 Design..114

6.7 RESULTS FOR EXPERIMENT 11 (PRE-CUEING)..114

6.8 DISCUSSION OF EXPERIMENTS 10 AND 11..118

6.8.1 Effects of attention on sampling..119

6.8.2 Effects of attention on internal noise..120

6.8.3 Relation to effects of attention on visual short-term memory..121

6.9 SUMMARY OF FINDINGS IN CHAPTER 6..121

7. SUMMARY AND FINAL REMARKS..123

8. REFERENCES...125

9. APPENDICES...141

APPENDIX 1: INDIVIDUAL ERROR DISTRIBUTION PLOTS............................141

APPENDIX 2: PREDICTIONS FROM A STRICT CAPACITY = 1 OBJECT MODEL...144

APPENDIX 3: PILOT EXPERIMENTS WITH CONSTANTLY VISIBLE TARGET MARKER CUES AND WITH NO NON-TARGETS..147

APPENDIX 4: INDIVIDUAL LAG PLOTS...152

APPENDIX 5: INDIVIDUAL LAG PLOTS FOR EXPERIMENTS 6 – 9................155

APPENDIX 6: MOT SIMULATIONS..159

APPENDIX 7: DUAL TASK PILOT EXPERIMENTS..163

APPENDIX 8: ANALYSES CONDUCTED ON DATA FROM EXPERIMENTS 10 AND 11 WITH CONSTRAINTS PLACED ON THRESHOLDS...168

APPENDIX 9: SIMULATED LAG PATTERNS..171

CHAPTER 1. GENERAL INTRODUCTION TO VISUAL ATTENTION

In everyday life we frequently refer to attention, or the process of "paying attention" to something, although it is not usually clear exactly what we mean by this, nor indeed what it is about attention that is useful to us. There is usually the notion that paying attention to something necessarily entails not paying attention to something else i.e. that the process of attention is one of selection. Inherent in this is the idea also of a limited capacity, after all we cannot pay attention to every detail of our environment. Thirdly, when we pay attention to something, we have the feeling that is allows us to perceive, think about or interact with it more effectively. It is these questions which are the focus of this thesis. How does attention help us to perform tasks effectively and what is the nature of the capacity limit of attention? In addition, does attention operate differently when we try to perform different tasks?

1.1 Models of general attention

Early general models of attention attempted to answer these questions in terms of a bottleneck in processing of sensory information. In Broadbent's filter model (1958), for instance, the semantic or 'meaning-related' content of information is only processed after attentional selection has taken place. This was supported by evidence that in dichotic listening tasks where two streams of speech are heard but only one actively attended, participants can often report the gender of the speaker in the unattended speech stream but not the content (Cherry, 1953). Conversely, in the late selection model (Deutsch and Deutsch, 1963), some semantic information is accessible for report even from apparently unattended sources, since all information reaches short-term memory before attentional selection takes place. It has been difficult to resolve this particular debate, however, because of the difficulty in establishing whether information is truly unattended. In any case, these early models of general attentional capacity contain the idea that there is a qualitative and/or quantitative limit of what can be attended at any one time.

1.2 Visual spatial attention

In the visual domain, much research was initially, and still is devoted to the spatial allocation of attention to visual stimuli. When we attend to a visual stimulus in the environment, it is usually natural to look at it i.e. to fixate it and thereby cause the image of the attended object to fall in the fovea. In 1980, Posner demonstrated empirically what seems apparent from our intuitions; that it is also possible to fixate one point, and attend elsewhere i.e. to use *covert* attention. In this thesis I shall use the term "attention" to refer to this form of processing which is independent of fixating the attended stimuli. In Posner's studies, observers fixated a central point, and a cue indicated that a target stimulus would appear either to the left or the right of fixation. The cue could either be valid (in that it correctly indicated where the target would appear), neutral (conveyed no information about where the target would appear) or invalid (indicating the opposite direction to where the target would subsequently appear. Observers were faster to react to the target when it had been validly cued than when the cue was neutral or invalid. They were also slower to react to the target when it had been invalidly cued than when the cue was valid or neutral. Posner attributed this result to the shifting of "covert" attention to the cued location. This time, in the visual domain, we see that attentional selection facilitates perception of selected objects, but is detrimental to perception of non-selected stimuli.

1:3 The necessity of attention for perception

More recently, new paradigms have been developed to investigate not what we *do*, but what we *do not* notice or perceive. Inattentional blindness and change blindness studies have shown the great extent to which visual information can be 'filtered out' by attention. Inattentional blindness is the phenomenon whereby observers completely fail to notice objects or events because their attention is focussed on a different task or stimulus. For instance, Mack and Rock (1998) report a striking case where observers are effectively 'blind' to a substantially salient superthreshold stimulus because of their concurrent engagement in another task. They presented observers with a cross and asked observers to make a length judgement about the two component arms. On some trials, they additionally presented a completely unexpected and task-irrelevant square. After observers had made their length judgement, all twelve observers could report something about the location of the square but three out of twelve reported that they were unaware that it was presented at all. Because their attention was directed elsewhere, they failed to consciously perceive the square. Interestingly, these results indicated that even those observers who did not consciously perceive the square appeared to have some degree of implicit processing, allowing them to report the quadrant correctly.

Simons and Chabris (1999) report a striking example of inattentional blindness. They asked observers to watch a video containing two teams passing basketballs between players and to count the number of passes made. During some trials, a person walked right across the scene wearing a gorilla suit! After the video had ended, observers gave their responses and then asked whether they had noticed anything unusual about the video, and if so, what it had been. Depending on the team they were asked to monitor, and the colour of the team's kit, between 42% and 83% noticed the person in the gorilla suit. Clearly, engagement in the pass-counting task and expectations about likely events had significantly affected which stimuli were accessible to conscious awareness: attending to black stimuli allowed the person in the gorilla suit to be noticed more frequently.

A related phenomenon to inattentional blindness, change blindness, is the phenomenon where surprisingly large changes can go undetected when attention is not focussed on the object or event of change (Intraub, 1997; Rensink, O'Regan & Clark, 1997; Simons & Levin, 1997). O'Regan, Rensink and Clark (1999) used pairs of photographs presented cyclically to examine how people detect changes in natural scenes. One of each pair had been modified (for instance, the location or colour of one item in the scene may have changed), and observers had to detect these changes. However, at the point at which the change occurred, the scene was presented along with a number of small areas likened to 'mud splashes'. Changes to objects rated as of 'central interest' were detected almost immediately, but changes rated as of 'marginal interest' were frequently not detected at all. As the mud splashes did not obscure the areas of change directly, O'Regan et al. (1999) interpreted the results as suggesting that they attracted attention away from the target areas by means of stimulus onsets involved (e.g. Jonides & Yantis, 1988; Yantis, 1993; Yantis & Johnson, 1990). That focussed attention is required to detect change is also indicated by the difference in detection rates between areas of the scene rated as being of different levels of interest to the observers.

Clearly, attention to visual stimuli has been shown empirically to display the intuitive notion of selection, or the facilitation of processing of attended stimuli at the expense of non-attended stimuli. This thesis will cover three questions about visual attention: what is it that attention does that is helpful to performing tasks, what is the nature of the capacity limit, and does attention operate differently when we try to perform different tasks?

1.4 Questions on attention in this thesis

In Chapters 2 - 4, we will investigate the nature of the capacity limit for simultaneously attending to the features of multiple objects and whether there is a serial component to apparently simultaneous processing. Any element of serial processing could in itself reflect a form of capacity limit, since it implies that any observed limits would be even stricter if serial processing were somehow prevented. That is to say, the presence of a serial component could inflate any apparent capacity limit based on the number of objects that are apparently attended simultaneously. Lastly, we ask if there a difference between the observed capacity limits for simultaneous attention for different features?

In Chapter 6, through the use of a very different paradigm, we will demonstrate that attention facilitates perception in a similar way for briefly presented stimuli as it did in Chapters 2-4 for attention to more temporally extended stimuli. In Chapter 6, the noise reduction capability of attention will be investigated in some detail, with reference to its similarity with the processes involved in the phenomenon of visual crowding. In Chapter

6, equivalent noise methods will be used to examine separately the effects of attention on sampling by the visual system of stimuli, and on internal noise reduction during encoding of the representation of target objects. The effects of two types of attentional manipulation; dual-tasks and spatial cueing, will be investigated separately, and compared with visual crowding, a phenomenon thought to affect internal noise levels.

CHAPTER 2. EXPERIMENTS 1 – 3: CONTINUOUS MONITORING OF THE FEATURES OF MULTIPLE OBJECTS

2.1 INTRODUCTION

2.1.1 Summary of literature covered in the Introduction

The nature of the capacity limit for attending to the features of multiple objects is an under-explored area. There are many studies examining performance in tracking the positions of multiple objects, but many suffer from inflation of set size effects due to decision noise, and none have measured positional errors directly. For non-positional features, there are very few studies examining the nature of a capacity limit on perception. Most address the issue indirectly through performance in short-term memory tasks, or through change detection, for instance, and these are tasks that may have different and additional demands other than perception of multiple objects.

2.1.2 Continuously monitoring positional features

The nature of the capacity limit for attending to the positions of multiple objects has been well investigated using variants of the multiple object tracking (MOT) paradigm originally developed by Pylyshyn and Storm (1988) (e.g. Alvarez & Cavanagh, 2005; Alvarez & Franconeri, 2005; Alvarez, Horowitz, Arsenio, DiMase & Wolfe, 2005; Keane & Pylyshyn, 2006; Pylyshyn, 1989; Scholl & Pylyshyn, 1999; Scholl, Pylyshyn & Feldman, 2001; vanMarle & Scholl, 2003; Yantis, 1992). Typically, in an MOT task, a number of identical objects appear, some are designated as targets to be tracked, and then all move randomly about the screen. This method is reminiscent of the classic carnival shell game, where the task is to keep track of one out of a number of identical moving objects. After the objects stop moving, observers report which were targets. Accuracy declines with number of targets, suggesting a capacity limit for tracking. The task provides a measure of sustained attention to the positions of multiple objects because observers must continuously update their representations of objects' positions. The MOT task measures the ability to keep track of which objects are targets and which are non-targets, and hence does not, measure the effect of number tracked on the size of the positional imprecision in localising targets. Although potentially the culprit for the limit on the number that can be tracked, these positional errors manifest only indirectly in the percent-correct measure.

2.1.3 The nature of the attentional limit

In the Pylyshyn and Storm (1988) paper, observers were unable to accurately perform the MOT task if there were more than four or five target objects, and Pylyshyn (1989) proposed that people have four or five pointers or "FINSTs" (or "fingers of instantiation") for tracking. In this sense, each object either has a FINST used to track it, or it does not, and tracking capacity is defined by this binary outcome. More recently, a number of researchers have argued against the idea of a fixed-number-of-objects limit on position tracking. These researchers (e.g. Alvarez & Cavanagh, 2005; Alvarez & Franconeri, 2007; Yantis, 1992) have shown a progressive decline in performance for each tracked object, as the number of tracked objects is increased. However, for the most part these studies are vulnerable to inflation of the set size effect by decision noise factors of the type discussed by Palmer (1995). If the internal representations of objects are noisy, then each will contribute its own chance of making an error on any judgement task. If the task requires decisions based on representations of multiple objects, then the likelihood of an error will increase with the number of objects, even in the absence of any decline in precision of each representation with set size.

2.1.4 Further evidence against an attentional limit based on a fixed number of objects

The only position monitoring study we are aware of which circumvents this type of problem is that of Alvarez and Cavanagh (2005, Experiments 1 & 2) although they only found a decline in processing of tracked objects when they were all presented to one hemisphere. Alvarez and Cavanagh were able to avoid decision noise by querying the representation of only one attended object at a time. Tripathy and colleagues successfully avoid decision noise influences and use an innovative task that tests the ability to detect a deviation in the trajectory of an object (e.g. Tripathy & Barrett, 2003a; Tripathy & Barrett, 2003b; Tripathy & Barrett, 2004). They find a very large effect of set size on performance. However, trajectory change detection requires detecting a change in an ongoing pattern of position change and may be performed through monitoring directions of motion or by a comparison process between different remembered positions of objects, for instance. Here, we investigate the capacity to perform the task of monitoring position itself, rather than change in change of position.

Franconeri, Alvarez and Enns (2007) have recently shown for spatial selection that the capacity limit depends on the degree of spatial precision required. They precued a variable number of locations before a visual search task for oriented lines, and found that the capacity limit (90% performance-rate defined) was dependent on the spatial density of the potential search locations. There are three reasons why their results cannot, however, be interpreted directly as reflecting the nature of the capacity limit for attending to multiple locations. Firstly, their manipulation of display density to estimate the precision of selection is also partly an estimate of the capacity to exclude nearby distracter objects from processing. We cannot know, from these experiments, the contribution that distracter exclusion made to the observed capacity limits. Secondly, their task included a 500 ms retention interval leaving open the possibility that their observed bottleneck resides at the level of visual short term memory. Thirdly, theirs was a visual search task for the orientations of objects at selected locations, and is hence a measure both of the spatial selection process and the subsequent encoding of objects' orientations. Here in Experiment 3 we more directly assess the processing of locations themselves without additional demands on encoding objects' additional features, without retention intervals and without confounding effects of distracter exclusion.

Similarly to Franconeri et al. (2007), Alvarez and Franconeri (2007) concluded that tracking performance in MOT studies is dependent on the spatial precision of attention, and that this can explain variable estimates of capacity limits when objects move at different speeds. They varied the minimum spacing between objects during an MOT task, and observed that capacity estimates were more dependent on minimum spacing for fast moving objects than for slow moving objects. Here in Experiment 3, we are able to assess this spatial precision directly, without needing to assume a direct causal link between spatial measures (such as minimum spacing) and object speeds.

2.1.5 Non-positional features

Few studies have investigated the capacity limit for attending to the non-positional features of multiple changing objects. While normally multiple objects maintain their separate identities by occupying distinct locations, Blaser, Pylyshyn, & Holcombe (2000) used objects that occupied the same location but maintained separate identities by having distinct features that nevertheless smoothly changed through "feature space". They asked observers to continuously monitor the spatial frequency, colour, and/or orientation of spatially superposed Gabor patches. Observers judged the direction of a "jump" in the Gabor patches' trajectories through feature space. Performance was so much better when monitoring two features of the same Gabor patch than for two features belonging to two different patches that the difference was best explained by a capacity limit of just a single object. However, it is unclear whether this result would generalise to objects in separate locations that might undergo less competition. Scholl, Pylyshyn and Franconeri (1999) had observers track multiple moving objects and found that encoding of objects' colours and shapes was quite poor, without the additional demand of monitoring multiple objects in a single location. Much of the difficulty however may have stemmed from the concurrent demand of position tracking: Saiki (2003) showed that feature change detection was very poor when tracking three objects. Saiki had observers detect a colour switch event where two of three moving objects exchanged colour with one another. The feature monitoring experiments presented here will begin by investigating non-positional feature monitoring by itself.

In a change detection task, Bahrami (2003) used the multiple object tracking paradigm to study attention to objects' features. During tracking, they introduced colour and shape changes to either tracked or untracked objects, either with or without a 'mudsplash' event similar to those used by O'Regan et al. (1999) which did not directly obscure the change event. They found that with the 'mudsplash', observers were likely to miss these changes even when they occurred in tracked objects. They interpreted this as evidence that attention had been captured by the mudspashes, thereby drawing it away from the transients produced by the feature changes. Furthermore, they concluded that directed attention to the feature change is necessary for change detection. Surprisingly, on trials without mudsplashes, detection rates for shape changes in tracked objects were under 60%, suggesting that object representations are relatively impoverished during MOT. Again, however, it is not clear how much the concurrent demand of position monitoring affected performance. In addition, change detection requires a comparison process between current and previous states of objects, and this may also have contributed to the poor levels of performance observed. Again, the feature monitoring experiments presented here will be free of these additional demands.

2.1.6 Non-positional features: visual short-term memory studies

In contrast to the paucity of studies investigating non-positional feature monitoring with continuously changing displays, there have been many studies using static displays to measure visual short-term memory (VSTM) for non-positional features. The findings of these studies are of interest for understanding the monitoring of objects because the capacity limits for VSTM and visual attention may be intimately related (e.g. Cowan, 2000).

In feature-change detection tasks, Luck and Vogel (1997) found an apparent four-object limit in a VSTM task, and more recently Alvarez and Cavanagh (2004) have observed a more progressive decrease in performance based on the number of objects for encoding, *and* their visual complexity. This has since been challenged, however, on the basis that the apparent effects of complexity can be removed when sample-test similarity is accounted for (Awh, Barton & Vogel, 2007). However, using a slightly different methodology and a more direct probe into representations of objects, Wilken and Ma (2004) probed memory for colour, orientation and spatial frequency. They asked observers to adjust the features of a test stimulus to match that which had been presented and observed a systematic decrease in precision as set size was increased progressively from two to eight objects.

Additionally, a number of studies have reported performance costs for change detection tasks involving multiple instances of the same feature dimension (e.g. Magnussen & Greenlee, 1997; Wheeler & Treisman, 2002; Xu, 2002 but see Luck & Vogel, 1997). In other words, they report a capacity limit on detecting changes in more than one value of a single feature; multiple different orientations, for instance. Despite this literature on VSTM and change detection tasks, the capacity for attending to the non-positional features of multiple objects remains unknown. When there is no memory component to the task, and when there is no comparison process required, it may be that a similar limit is observed for multiple values of the same feature, but this has yet to be determined empirically. A continuous monitoring task might yield a quite different result from these VSTM and change-detection tasks as it does not include the need to maintain memory over an interval. Continuous monitoring tasks might tap into the ongoing experience of the visual field that seems to be immediately available. Although it has been suggested that this amounts to an illusion of seeing (e.g. Blackmore, Brelstaff, Nelson & Troscianko, 1995; O'Regan, 1992) and indeed results from change-detection tasks are consistent with this, on the other hand the experience of a rich visual field might be accompanied by an ability to accurately report the features of any of several objects if they are immediately queried upon disappearance. Although the classic conception of a high-fidelity iconic memory (Sperling, 1960) might suggest no decrement with number of items monitored, this is not what we will find.

2.1.7 Parallel and serial processing

If there is a severe capacity limit in continuous processing of the features of several objects, our feeling of a rich visual field might partially reflect accumulation of information taken from previous moments, such as attentional samples. Such a process would result in our representation of some objects lagging behind the present, as information has not been sampled from those objects for a while. In the worst-case scenario of an object monitoring capacity limit of only one object, accompanied by serial processing where attention visits successive objects to sample one at a time, we would expect an increase in perceptual lag with each addition to set size because on average, the greater the number of objects tracked, the longer in the past any object will have last been visited by attention. However, it appears that such potential perceptual lags have not been isolated in psychophysical data.

Many studies have attempted to measure the speed with which attention can be transferred between objects. These studies have typically reported values in the range of several hundred milliseconds (e.g. Carlson, Hogendoorn & Verstraten, 2006; Duncan, Ward & Shapiro, 1994; Horowitz, Holcombe, Wolfe, Arsenio & DiMase, 2004; Reeves & Sperling, 1986; Ward, Duncan & Shapiro, 1996; Weichselgartner & Sperling, 1987). For instance, Carlson, Hogendoorn and Verstraten (2006) measured the speed of attention shifts using a 'clock monitoring' paradigm. They presented ten 'clock-like' stimuli with arms rotating at constant speeds but each with an independently determined orientation at the start of the trial. In the baseline condition, observers knew from the start which 'clock' would be cued, after which they should report the 'time' on that clock. In two further conditions, either a peripheral or central cue indicated which clock should be reported and observers reported the perceived reading at the time of the cue. They measured a 140 ms latency for peripheral cueing and a 240 latency for central cueing.

However, these studies have had to induce serial processing in order to measure attention shifts. For instance, Duncan et al. (1994) presented two to-be-attended stimuli in succession in order to induce a switch of attention. It is not at all clear whether these relatively slow rates of attention switches would have been observed in a task where simultaneous attention to multiple objects was encouraged or even required. Indeed, in visual search, where there is no explicit requirement to process objects serially, some suggest that attention switches among objects at a much faster rate (Wolfe, Alvarez, & Horowitz, 2000).

Here we present a task where truly parallel, simultaneous processing of multiple objects is encouraged, and indeed where only simultaneous processing will produce the most accurate responses. Because of the continuous nature of our task, we are able to introduce a new analysis that is sensitive to the presence of perceptual lags.

2.2 METHOD FOR EXPERIMENTS 1 AND 2: ORIENTATION AND SPATIAL PERIOD TRACKING

In Experiment 1, observers had to continuously monitor the changing orientations of a variable number of target Gabor patches. On all trials, five Gabors were displayed equidistant from a central fixation point (see Figure 1). At the start of each trial, either one, two or four of the possible Gabor locations were marked by black circles to indicate the locations at which targets would appear. All Gabors then independently underwent a period of smooth and quasi-random change in their orientations, and observed kept track of the orientations of target Gabors. After a quasi-random interval, all Gabors disappeared, and the location of one of the target Gabors was queried with a white circular marker. Observers used a keypress to adjust the orientation of a pair of black markers to match the last orientation of the queried target. Feedback was provided in the form of the queried Gabor with its final orientation.

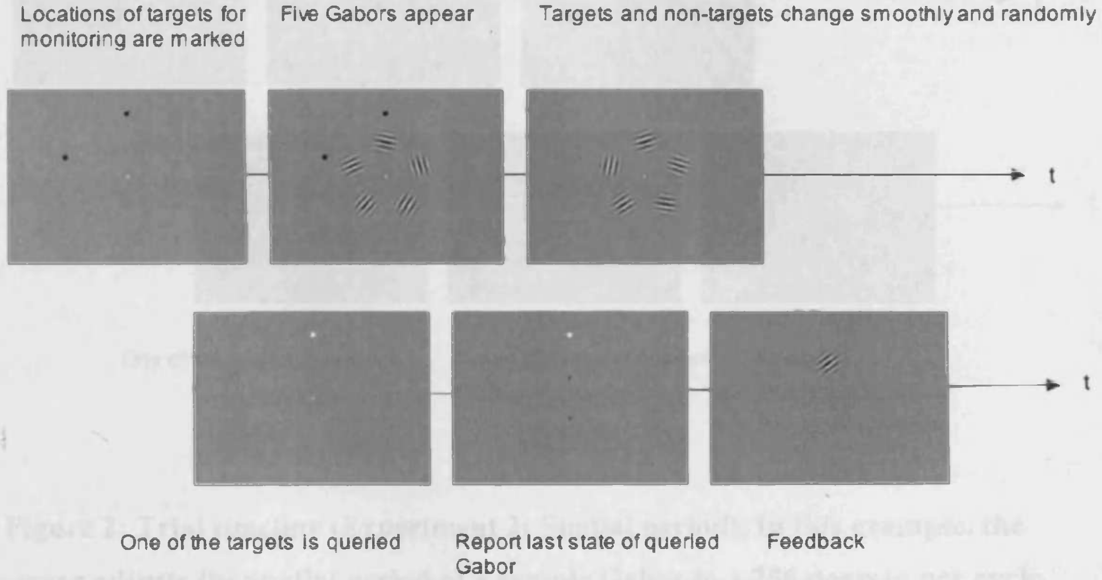

Figure 1: Trial timeline (Experiment 1: Orientation). Observers adjust the orientation of a pair of markers to match last state of the queried Gabor. In this example, the observer reports an orientation of 0 degrees. Given that the last

orientation of the queried Gabor was 50 degrees, this represents an error of -50 degrees.

The method for Experiment 2 was very similar to that for Experiment 1 with the following exceptions. The five Gabors now remained at a constant orientation, and instead changed independently in a smooth and quasi-random manner in their spatial periods. Observers tracked the spatial periods of the target Gabors. At the end of the trial, observers used a keypress to adjust the spatial period of a sample Gabor patch until it matched the last spatial period of the queried Gabor (see Figure 2). Feedback was provided in the form of the queried Gabor with its final spatial period.

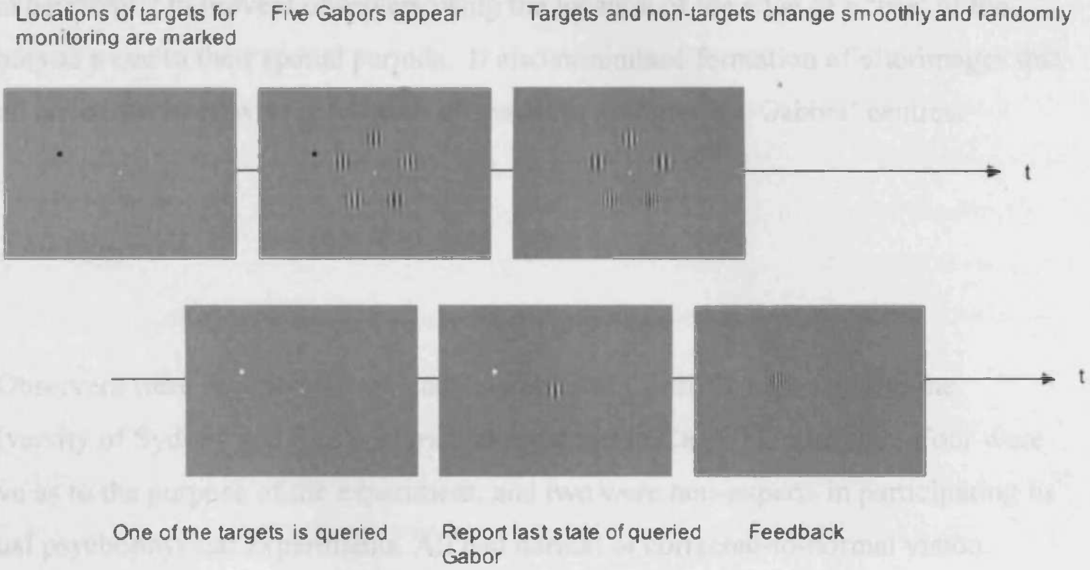

Figure 2: Trial timeline (Experiment 2: Spatial period). In this example, the observer adjusts the spatial period of a sample Gabor to 1.256 degrees per cycle. The last spatial period of the queried Gabor was 0.722 degrees per cycle, and this represents an error of +0.534 degrees per cycle.

A computer programme was written in Python using the VisionEgg library (http://www.visionegg.org) and displayed an array of sinusoidal Gabor luminance gratings against a mid-grey background on a 16 inch CRT screen refreshing at 85 Hertz. Observers viewed the display in a dimly lit room from a distance of 0.4m. The luminance of Gabors varied from 0.02 (trough) to 120.00 (peak) candelas per m^2. In Experiment 1 (orientation monitoring) Gabors had a constant spatial period of 1.062 degrees per cycle but a variable orientation. In Experiment 2 (spatial period monitoring) Gabors had a variable spatial period but fixed orientation of 0 degrees (vertical). The Gaussian envelope that windowed the Gabor patches' amplitudes had sigma = 1.139 degrees of visual angle. In Experiment 1, phase was such that the centres of Gabors had the maximum luminance defined by the sinusoidal function. In spatial period monitoring (E2), the phase of each Gabor was randomised from trial to trial. Phase was randomised in Experiment 2 to prevent observers using the location of the edge of a 'bar' of the Gabors as a cue to their spatial periods. It also minimised formation of afterimages that could have interfered with perception of spatial period near the Gabors' centres.

2.2.1 Observers

Observers were two psychology staff members at Cardiff University and the University of Sydney and four postgraduate students at Cardiff University. Four were naïve as to the purpose of the experiment, and two were non-experts in participating in visual psychophysical experiments. All had normal or corrected-to-normal vision.

2.2.2 Procedure

Observers were given practice trials until they felt comfortable with the experiment (usually less than 10 trials).

At the start of each trial either one, two or four black circular markers with a radius of 0.48 degrees of visual angle indicated the positions at which targets for monitoring would appear. These were presented peripheral to the future locations of Gabors, 11.94 degrees eccentric from the central fixation point.

After 2350 ms, five Gabors appeared in addition to the target markers. Gabors were presented equidistant from a central fixation point at the vertices of an imaginary pentagon such that their centres were always 6.79 degrees eccentric from the fixation point. The spacing between adjacent Gabors was 7.985 degrees. This configuration was chosen to avoid crowding, as crowding should not occur when spacing between objects is larger than half of their eccentricity (Pelli, Palomares and Majaj, 2004).

On any given trial, there was an equal chance of one, two or four Gabors being marked as targets for monitoring. Which of the five made up this target set was entirely random on each trial. Target markers remained on screen for the first 1180 ms of Gabor motion.

Each object stayed at the same spatial location throughout the trial, but varied in either orientation (E1) or spatial period (E2) according to a quasi-random trajectory around a 'feature space'. All Gabors moved around this feature space at all times according to an algorithm described in the "Trajectories" section (section 2.2.4). At a point randomly varying between 5350 ms and 10350 ms after the start of the trial (3000 – 8000 ms since appearance of Gabors), all Gabors disappeared. Immediately after the disappearance of the Gabors (12ms later i.e. on the next screen refresh), the location of one of the target Gabors was post-cued, and observers attempted to report the last orientation or spatial period of that object the moment it disappeared. The post-cue was identical to the target marker that had previously occupied that location except that it was white instead of black.

Observers reported the last orientation (E1) or spatial period (E2) of the queried object by adjusting a sample presented at the centre of the screen. As soon as observers started to make their adjustment using a keypress (on the first screen refresh after a keypress was detected), the sample stimulus appeared. In Experiment 1 it was a pair of adjustable markers whose orientation was controlled by keypresses (see Figure 1). In Experiment 2 it was an adjustable test Gabor whose spatial period was controlled by keypresses (see Figure 2). We delayed the appearance of the sample stimulus until after the observer made their first keypress to avoid any potential interference of the sample stimulus on the effort of the observer to recall the feature value.

For the orientation experiment, we chose a central sample because, if we had presented the sample at the same location as the queried Gabor, this might yield an apparent motion cue between the actual to-be-reported stimulus and the sample. This might allow observers to report the last orientation of the Gabor as a result of the motion signal and not as a result of them attending to that Gabor at the moment that it disappeared.

As shown in Figure 1, the pair of black markers were initially oriented at 0 degrees (vertically above and below the location previously occupied by the fixation point and each 4.110 degrees away from the centre of the screen). We chose a constant starting orientation for the markers for two reasons: to allow observers to become practiced at adjusting orientation from a given starting orientation, and secondly, since the ending orientation of the Gabor was random, to ensure that the degree and direction of adjustment required was also random on each trial. The orientation of the two markers could be adjusted with a keypress to that of the post-cued object just before it disappeared.

For the spatial period experiment, the sample stimulus was a sample Gabor patch with completely randomised phase at the central fixation point, identical to the previous five Gabors and with a starting spatial period of 0.95 degrees per cycle. Again, it was necessary to present the sample at a location other than that of the queried Gabor to prevent any motion signals being produced which could have allowed observers to access the previous spatial period of the queried Gabor in the absence of attention. We used the same marker as was presented at the start of the trial to indicate which of the Gabors was post-cued, except that it was white instead of black. Observers then adjusted the spatial period of this sample Gabor patch with keypresses until they felt it matched that of the queried Gabor at the time it disappeared.

At the end of each trial in both experiments, feedback was presented in the form of a static display containing only the queried Gabor in its last state before it disappeared.

2.2.3 Design

Observers completed three blocks of 105 trials in both experiments. Within each block there were 35 trials each for one, two and four targets for monitoring. In total, this yields 105 trials for each condition of one, two or four targets. On each trial, the number of targets for monitoring was selected randomly until at the end of each block each condition had been run 35 times.

2.2.4 Trajectories of Gabors through orientation or spatial period

The orientation (E1) or spatial period (E2) of each Gabor stimulus over time corresponded to a random trajectory through feature space and was generated by the following algorithm. Every 20 frames, corresponding to 235 ms, the acceleration of the Gabor through feature space would be randomly reassigned to positive or negative. If the Gabor had been changing slowly, the two possible accelerations were larger than if the Gabor had been changing quickly. This was to prevent the features of any particular Gabor remaining relatively constant for a prolonged period that could result in an afterimage forming. These changes in acceleration were usually not salient to observers, consistent with humans' low sensitivity to acceleration (Werkhoven, Snippe & Toet, 1992).

2:2.5 Trajectory parameters. Experiment 1: Orientation

Objects' starting orientations were set to random values. Starting angular velocities were set randomly and independently between -0.042 degrees per ms and 0.042 degrees per ms, excluding absolute values less than 0.0085 degrees per ms, such that no Gabors appeared to be stationary when first presented. Starting angular accelerations were randomly and independently chosen as either -2.17×10^{-4} degrees per ms^2 or 2.17×10^{-4}

degrees per ms^2. Every 235 ms, the angular acceleration was again set randomly to either -7.23×10^{-5} deg per ms^2 or 7.23×10^{-5} deg per ms^2. If the angular velocity was below 0.017 degrees per ms, then the absolute value of the acceleration was increased from 7.23×10^{-5} deg per ms^2 to 2.17×10^{-4} deg per ms^2. A maximum absolute value of angular velocity was set at 0.26 degrees per ms. If any Gabor reached this value, the direction of its acceleration was reversed such that it tended back towards slower velocities.

2.2.6 Trajectory parameters. Experiment 2: Spatial period

For this experiment, the starting spatial period of each Gabor was set independently to a random value between 0.7 and 1.2 degrees per cycle. We parameterised the changes in terms of degrees per cycle (dpc) or "bar" width rather than cycles per degree because it led to the changes appearing more uniform across the range of spatial periods. The velocity through spatial period space, or rate of change of degrees per cycle, was set at the start of each trial randomly and independently for each Gabor between +/- 0.000425 dpc per ms, ensuring that no Gabor had an absolute velocity below 8.5×10^{-5} dpc per ms. The starting accelerations were again randomly chosen each to be either 3.61×10^{-7} dpc per ms^2. Every 235 ms, the acceleration of each Gabor was re-set to either $+/-3.61 \times 10^{-7}$ dpc per ms^2. If the velocity was smaller than an absolute value of 0.000425 dpc per ms, the absolute value of the acceleration was increased to 3.61×10^{-6} dpc per ms^2.

During the trial, the maximum spatial period was set to 0.4 dpc and the minimum to 1.5 dpc. If the maximum or minimum values were reached, the sign of the velocity was changed such that spatial periods moved back towards the middle of the range of possible values. If the velocity reached a maximum absolute value of 0.00425 dpc per ms, the direction of acceleration would be reversed such that the velocity tended back towards lower values.

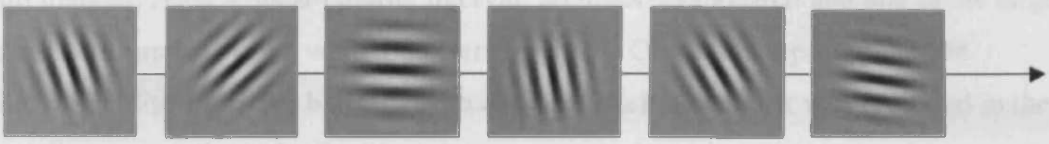

Figure 3: Example frames from the stimulus sequence for one particular Gabor, taken at 400 ms intervals (Experiment 1: Orientation). In the actual stimulus, the orientation changes in a smooth and quasi-random manner. The rate and direction of change are frequently varied, as is the acceleration through orientation values.

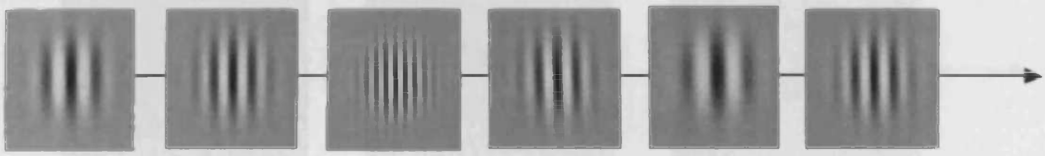

Figure 4: Example frames from the stimulus sequence for one particular Gabor, taken at 400 ms intervals (Experiment 2: Spatial period). In the actual stimulus, spatial period changes in a smooth and quasi-random manner. The rate and direction of change, as well as the acceleration through spatial period values are frequently varied. Spatial period values are limited to range between 0.4 and 1.5 degrees per cycle.

2.3 EXPERIMENT 3: POSITION TRACKING

In Experiment 3, eight black discs appeared on all trials at locations equidistant from a central fixation point. Each disc appeared in its own “ballcage” which limited its subsequent movement. At the start of each trial, either one, two, four, six or seven of these discs flashed for a brief period to indicate their status as targets for tracking. All discs then independently moved around inside their ballcages in a smooth and quasi-random manner. After a quasi-random interval, all discs disappeared and one of the target ballcages was queried with a white line (see Figure 5). Observers reported the last position of the disc from this ballcage with a mouse click. Feedback was provided in the form of the queried disc at its final location.

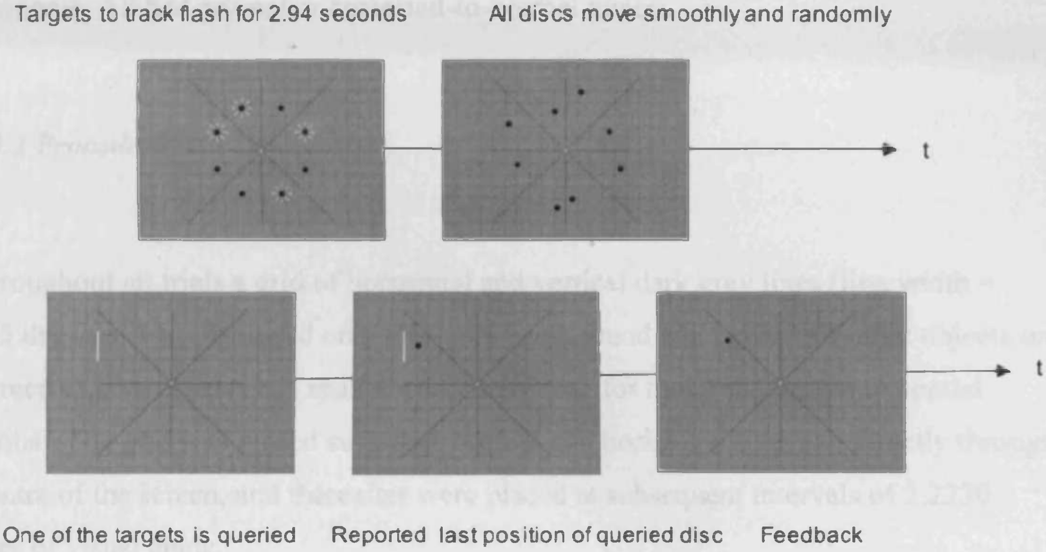

Figure 5: Trial timeline (Experiment 3: Position). Observers move a sample disc using the mouse to match the last position of the queried disc. In the example depicted in Figure 5, the observer reports a horizontal position relative to the bottom left corner of the screen of 11.71 degrees and a vertical position of 22.90 degrees. Given that the last horizontal and vertical positions of the queried disc were 12.98 and 23.65 degrees respectively, this represents an error of magnitude 1.47 degrees, at an angle of 239.76 degrees clockwise from vertical (position reported was below and to the left of the actual position).

The general method for Experiment 3 was the same as for Experiments 1 and 2, apart from the following differences.

2.3.1 Observers

Observers were one psychology staff member at the University of Sydney and six postgraduate students at Cardiff University. Three were naïve as to the purpose of the experiment, and two were non-experts in participating in visual psychophysical experiments. All had normal or corrected-to-normal vision.

2.3.2 Procedure

Throughout all trials a grid of horizontal and vertical dark grey lines (line width = 0.0445 degrees) was displayed on a mid-grey background and behind all other objects on the screen to give observers a spatial reference frame for their judgements of spatial locations. The grid was placed such that vertical and horizontal lines ran directly through the centre of the screen, and thereafter were placed at subsequent intervals of 2.2230 degrees of visual angle.

Eight triangular 'cages' were arranged around the central fixation point and limited the range of motion of the discs (Figure 6). Eight radial lines (luminance 0.02 candelas per m^2) separated these cages from one another and were centred on the fixation point. The outermost, eccentric edge of each cage was not marked on the screen but together the cages were circumscribed by an imaginary square subtending 25.0 degrees by 25.0 degrees in total, with the fixation point at its centre.

On each trial, eight black discs appeared with radius = 0.37 degrees. One appeared in each cage, beginning at an eccentricity of 8.83 degrees, situated at a point equidistant from the two closest radial cage walls. For the first 2.94 seconds after their appearance, either one, two, four, six or seven of the eight discs flashed repeatedly black and white to indicate their status as targets for tracking. On any given trial, there was an equal chance of there being one, two, four, six or seven targets, and which of the eight made up this target set was also entirely random on each trial. After the initial cueing period, all discs moved around their cages according to the trajectories described below.

At a point randomly varying between 3 and 8 seconds after the start of the discs' motion (to prevent observers from only paying attention just before the end of the trial), all discs disappeared, and the last location of one of the target discs was queried. Which of the target discs was queried was entirely random. The queried cage was indicated with a white marker line subtending 0.22 by 5.54 degrees that appeared at the outer edge of the cage. Observers attempted to report the last spatial location of that object the moment before it disappeared by using a mouse. As soon as observers started to make their adjustment by moving the mouse towards the perceived last location, the sample disc appeared at the centre of the screen. This was a disc identical to the eight discs previously on screen. We delayed the appearance of the sample stimulus until after the adjustment had started to avoid any potential interference in memory or perception between the location of the sample disc and the location reported by the observer. As soon as observers moved the mouse, the sample disc moved in the same direction as the mouse. After observers clicked the mouse to indicate they had completed adjustment of the sample, feedback was immediately presented by displaying the queried disc in its actual last position before it disappeared.

As the moving discs were confined to their own triangular cages, there was little chance of confusing a target with another disc. This improves the reliability of our probe of the observers' representations of the targets.

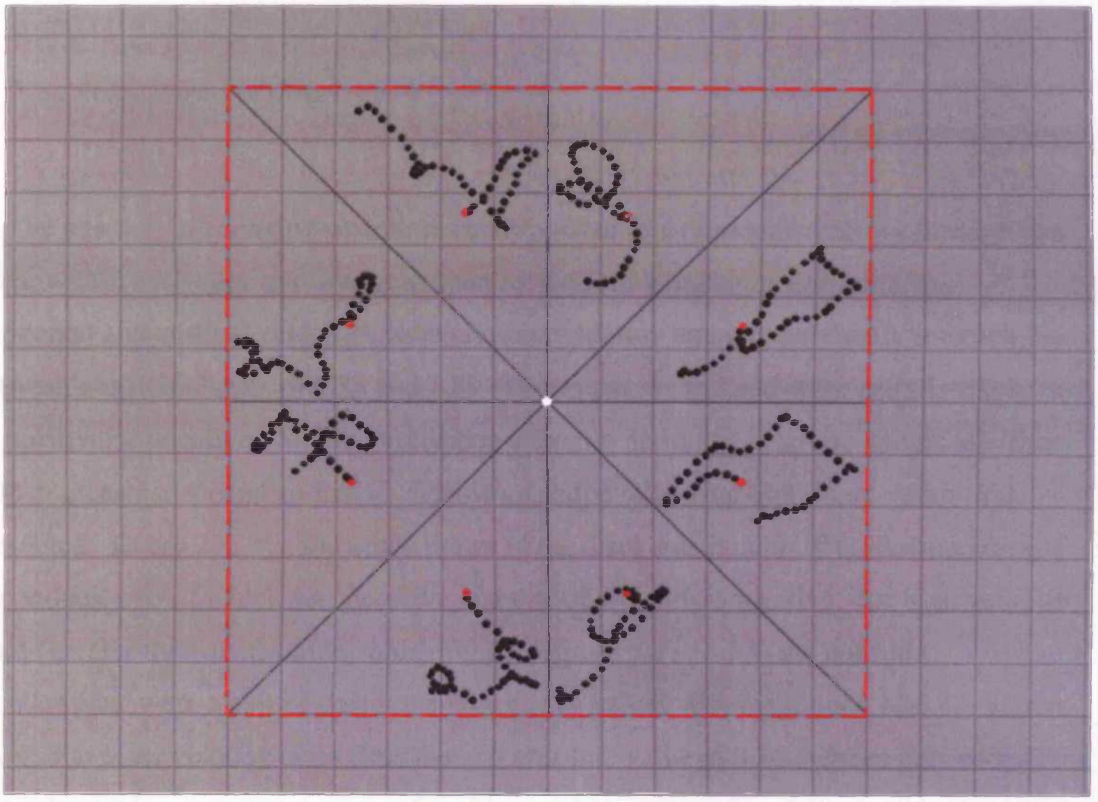

Figure 6: Trajectories of the eight discs in a typical trial of Experiment 3, position tracking. The trial depicted here lasted 5.5 sec. Plotted points mark the location of the centre of each disc every 120ms. Red points indicate starting positions. Red lines were not shown in the experiment but here indicate the outer edges of triangular cages.

2.3.3 Design

Observers completed three blocks of 175 trials. Every trial contained eight moving discs. Within each block there were 35 trials for each condition: tracking one, two, four, six and seven targets. In total, this yields 105 trials for each number of targets for tracking: one, two, four, six or seven. In each block, the number of targets for tracking was selected randomly until in each block each condition had been run 35 times.

2.3.4 Trajectories of discs

The position of each disc over time corresponded to a random trajectory through the space within each cage and was generated by the following algorithm. Starting horizontal and vertical velocities were chosen randomly and independently for each disc between absolute values of 0.38 and 1.89 degrees per second and were equally likely to be positive or negative. Starting velocities were not permitted absolute values less than 0.38 degrees per second so that no Gabors appeared to be near-stationary when first presented. Every 235 ms, the acceleration of the discs was randomly reallocated one of two values, one of which was negative and one of which was positive. For high velocity discs (greater than an absolute value of 0.76 degrees per second) the possible accelerations were a pair of positive and negative values with magnitude of 6.43 degrees per s^2. For slow velocity discs (less than an absolute value of 0.76 degrees per second) the possible accelerations were a pair of positive and negative values with greater magnitude of 16.07 degrees per s^2. This was to prevent any particular disc remaining relatively still for a prolonged period that could result in an afterimage forming. The velocity and current position values were then calculated for each frame. A maximum absolute value of velocity was set at 3.78 degrees per s with a mean velocity of 3.22 degrees per s. If any velocity reached this value, the acceleration direction was reversed, causing it to tend back towards slower velocities. Discs were bounded within their cages by checking for locations within 0.44 degrees of the edges of the cages. If the location of a disc reached this value, the direction of the velocity was changed, causing it to apparently reflect off the boundary.

2.4 RESULTS AND DISCUSSION FOR EXPERIMENTS 1-3

To examine precision of the observers' representation of the features reported, we plotted error histograms for orientation (Experiment 1), spatial period (Experiment 2) and position (Experiment 3). These are shown in figures 7 - 9.

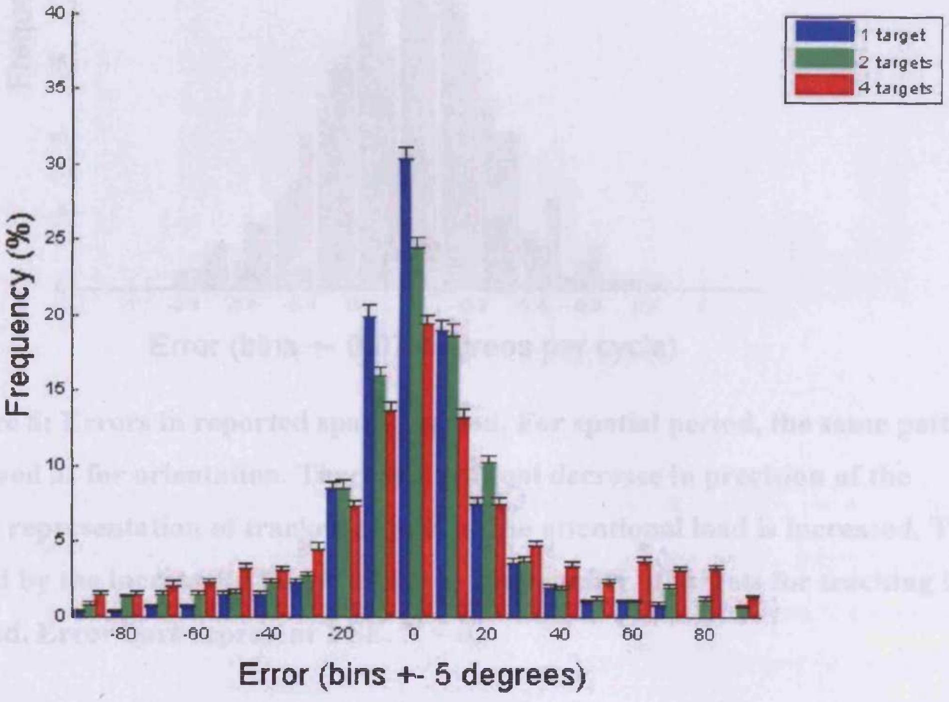

Figure 7: Errors in reported orientation. Error distributions for the three attentional load conditions show a gradual decrease in precision with increases in the number of objects tracked. The proportion of near-zero errors decreases with each addition to attentional load, and this is true even comparing tracking one and tracking two Gabors. Error bars represent 1 SE. N = 6.

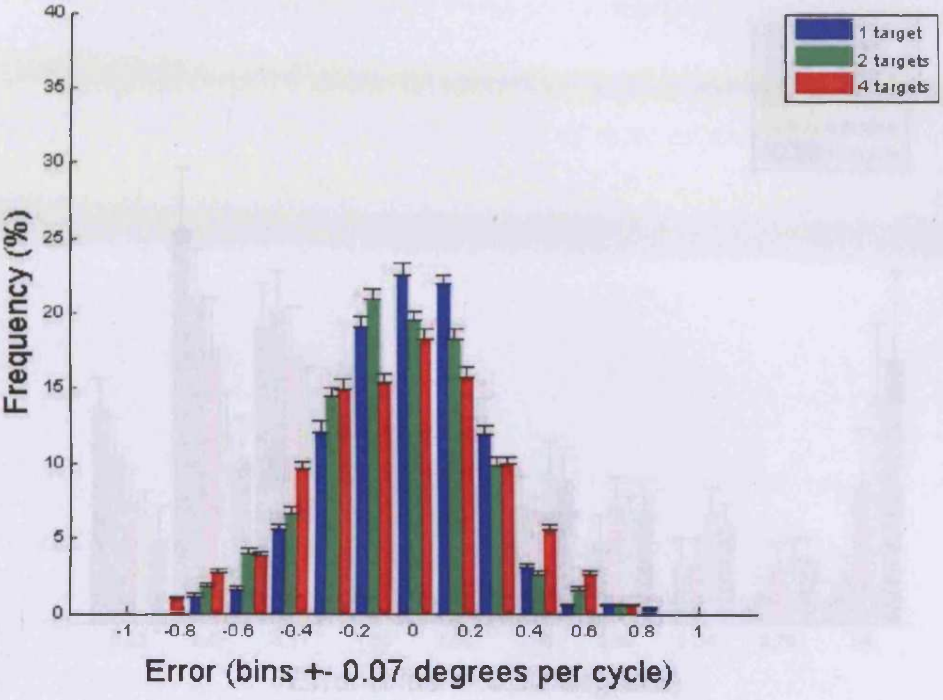

Figure 8: Errors in reported spatial period. For spatial period, the same pattern is observed as for orientation. There is a gradual decrease in precision of the internal representation of tracked objects as the attentional load is increased. This is reflected by the increase in larger errors as the number of targets for tracking is increased. Error bars represent 1 SE. N = 6.

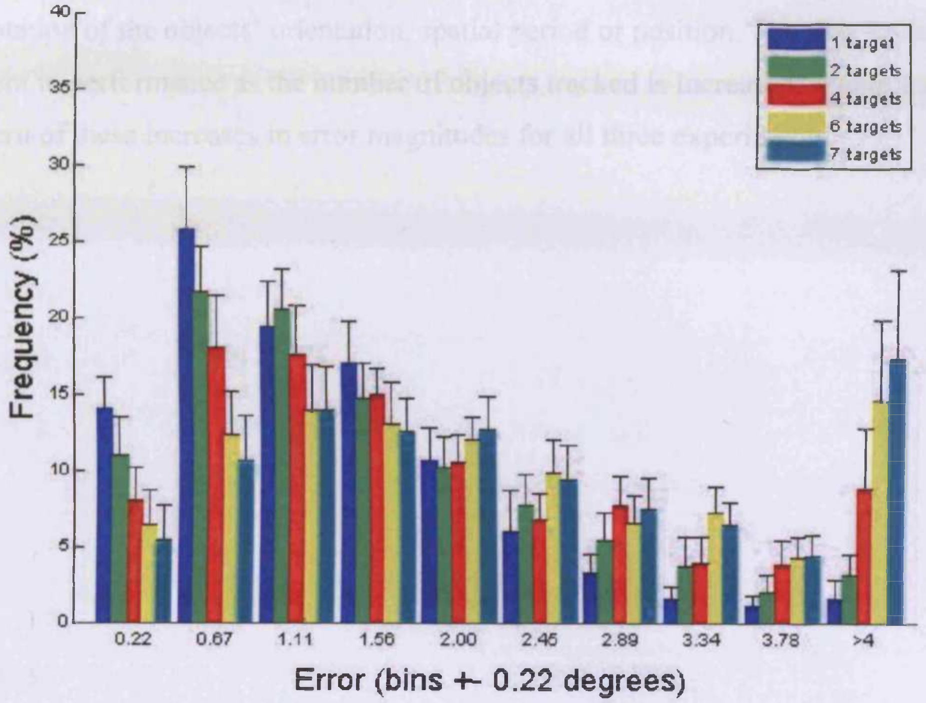

Figure 9: Errors in reported position for tracking one, two, four, six and seven objects. Error magnitudes are grouped into bins 0.45 degrees wide. For instance, the first bin (for all five set sizes) represents the frequency of errors between 0 and 0.45 degrees. Apart from the small spatial bias at small error magnitudes (described in the text), the distribution of errors for position tracking shows the same pattern as for orientation and spatial period tracking. Additions to the number of targets for tracking are associated with decreases in the number of near-zero errors, and increases in errors of greater magnitude. Error bars represent 1 SE. N = 7. Error magnitudes are measured in terms of the distance between the reported and actual last position of the queried object. On one of all the trials, the error magnitude was greater than the maximum possible response still within the correct cage. This data point was recoded as the same value as the otherwise most extreme error magnitude (13.33 degrees).

Broader distributions indicate larger average errors, or in other words, a less precise representation of the objects' orientation, spatial period or position. The data reveal a decrement in performance as the number of objects tracked is increased. Table 1 shows the pattern of these increases in error magnitudes for all three experiments.

E1 (orientation)	**Error variance (deg^2)**	**Mean error magnitude (deg)**	**% responses <= 18.4 deg (absolute values)**
Tracking one	**432**	**14.1**	**75.0**
Tracking two	**804**	**19.2**	**67.3**
Tracking four	**1360**	**27.0**	**51.6**
E2 (spatial period)	**Error variance (dpc^2)**	**Mean error magnitude (dpc)**	**% responses <= 0.284 dpc (absolute values)**
Tracking one	**0.0615**	**0.196**	**75.0**
Tracking two	**0.0699**	**0.215**	**72.7**
Tracking four	**0.0953**	**0.253**	**62.2**
E3 (position)	**Error variance (deg visual angle^2)**	**Mean error magnitude (deg)**	**% responses <= 1.72 deg (absolute values)**
Tracking one	**0.84**	**1.31**	**75.0**
Tracking two	**1.69**	**1.57**	**66.3**
Tracking four	**2.51**	**1.94**	**56.5**
Tracking six	**2.81**	**2.32**	**44.1**
Tracking seven	**3.23**	**2.46**	**41.2**

Table 1: Error variances, mean magnitudes and frequencies of small absolute errors. The trend towards less precise representations of tracked objects with greater attentional load is illustrated by the variance of errors, the mean error magnitudes and the frequencies of small error magnitudes. Small absolute error

magnitudes were arbitrarily defined as those for which the tracking one condition yielded 75% of responses.

Levene's tests for equality of variance (Levene, 1960) confirm that the distribution of errors increases with load, with the exception of the difference between the one and two object conditions for spatial period, and between the six and seven object conditions for position. Orientation monitoring (E1): difference between tracking one and tracking two Levene statistic = 24.095, $p < 0.01$, difference between tracking two and tracking four Levene statistic = 36.684, $p < 0.01$[1]. Spatial period monitoring (E2): difference between tracking one and tracking two Levene statistic = 2.759, $p = 0.097$, difference between tracking two and tracking four Levene statistic = 16.642, $p < 0.01$. Position monitoring (E3): difference between tracking one and tracking two Levene statistic = 22.31, $p<0.01$, between tracking two and tracking four Levene statistic = 24.37, $p<0.01$, between tracking four and tracking six Levene statistic = 6.60, $p<0.05$, but between tracking six and tracking seven Levene statistic = 1.56, $p = 0.21$.

2.4.1 Summary of set size effects on precision

This monotonic increase in the noise of representations of objects, as measured by the variance of errors, is not consistent with a fixed capacity model of attention. This would

---

[1] Because of the unique characteristics of the orientation experiment, that chance levels of performance would produce a flat, uniform distribution of error magnitudes, we were able to assess whether some errors resulted from erroneously reporting the last value of a non-queried target. For those trials when observers monitored the orientations of two Gabors, we analysed the difference between responses and the last value of the non-queried Gabor. There was no suggestion from the data that the distribution of errors differed from a uniform distribution of error magnitudes. In other words, there was no evidence of erroneously reporting the "wrong target".

predict no increase in standard deviations until the object limit was reached as observers only needed to make a decision about one object in each trial. The results show that attention cannot be distributed over multiple items without some decrease in precision either because of limited-capacity, parallel processing or the presence of a serial component.[2]

We also plotted data for individual observers shown in Appendix 1. The same broad pattern of a decrease in precision is observed across the data, although different observers had different average accuracy.

2.4.2 Spatial bias in Experiment 3

In the position-monitoring experiment (E3) we observed a small bias in responses such that the error distribution did not peak at zero: on average, there was a tendency to report values in roughly the same angular direction as the queried object from the centre, but a different distance, and this was more true in the upper hemifield. The overall tendency was to report values slightly further towards the top of the screen and further out from the actual last position of the queried disc. There was an overall mean signed error of 0.17 degrees upwards of last position.

For the left hemifield discs, the tendency was to report positions 0.07degrees further to the left than the last actual position, and correspondingly for the right hemifield there was a tendency to report positions the same distance further to right.

---

[2] The decline in performance with set size is in fact probably even worse than might be interpreted from the data, since performance could be contributed to by guessing strategies. This possibility and some preliminary modelling of the implications for orientation monitoring are considered in Appendix 2.

2.4.3 Distracter effects and effect of target marker duration

For orientation and spatial period monitoring, we also ran pilot studies where the target marker cues stayed on throughout the presentation of Gabors, to eliminate the possibility that the observed set size effect might be due to forgetting which Gabors were the targets. We also ran a pilot where only target Gabors appeared on screen to eliminate the possibility that the observed set size effects were due to exogenous attention capture from the non-targets. The results, presented in Appendix 3, indicate little or no effect of these factors.

2.4.4 Perceptual lag analysis

Errors represented up to this point have represented the difference between the reported value and the corresponding feature state of the queried object in the last frame before its disappearance. But if there were a lag in perception, then responses may be more similar to previous than to final feature states of the queried object. To assess this, we calculated mean absolute errors not just between responses and last states of the queried object, but with states of the object on every frame during the 40 last frames (470 ms) before disappearance of the objects. On these plots, the minimum error should occur at the point of mean perceptual lag.

Consider for instance the data points shown in blue in Figure 10. The leftmost point represents the mean of all error magnitudes that were previously shown in the histogram in Figure 7 for tracking the orientation of one object. This point represents the mean magnitude of the difference between the actual last orientation of the queried object, and the reported last orientation of the queried object. The second point along the blue curve in Figure 10 represents the mean magnitude of the difference between the reported last orientation of the queried object and the actual orientation 12 ms before it disappeared. Comparisons are made in the same way for the difference between reported last orientations and the orientations at every frame back through time before offset of the stimuli. Perceptual lags show up in this analysis as minima on the curve of mean error magnitudes, since these reflect the times at which the queried object was most similar to the reported values.

Perceptual lags associated with tracking orientation (shown in Figure 10) increased with the number of objects tracked: no lag when tracking one Gabor, 10 ms when tracking two, and 40 ms when tracking four Gabors (errors bars represent standard errors: +- 4 ms for tracking one, +- 6 ms for tracking two and +- 21 ms for tracking four). Perceptual lags associated with tracking spatial periods (shown in Figure 11) are much greater: 140 ms when tracking one Gabor, 210 ms when tracking two, and 250 ms when tracking four Gabors (errors bars represent standard errors: +- 16 ms for tracking one, +- 23 ms for tracking two and +- 35 ms for tracking four). For tracking positions, measured perceptual lags (shown in Figure 12) were 40 ms for tracking one object, 50 ms for tracking two, 90 ms for tracking four, 90 ms for tracking six, and 130 ms for tracking seven objects (errors bars represent standard errors: +- 32 ms for tracking one, +- 28 ms for tracking two, +- 26 ms for tracking four, +- 28 ms for tracking six and +- 129 ms for tracking seven).

In Experiment 1 (orientation), the correlation between the lag time and the number of objects tracked was highly significant (correlation coefficent = 0.551, N = 18, p (1 tailed) < 0.01). This was similarly the case for spatial period tracking in Experiment 2 (correlation coefficient = 0.609, N = 18, p (1 tailed) < 0.01). In Experiment 3 (position) the correlation was nearly significant (correlation coefficient = 0.276, N = 35, p (1 tailed) = 0.054).

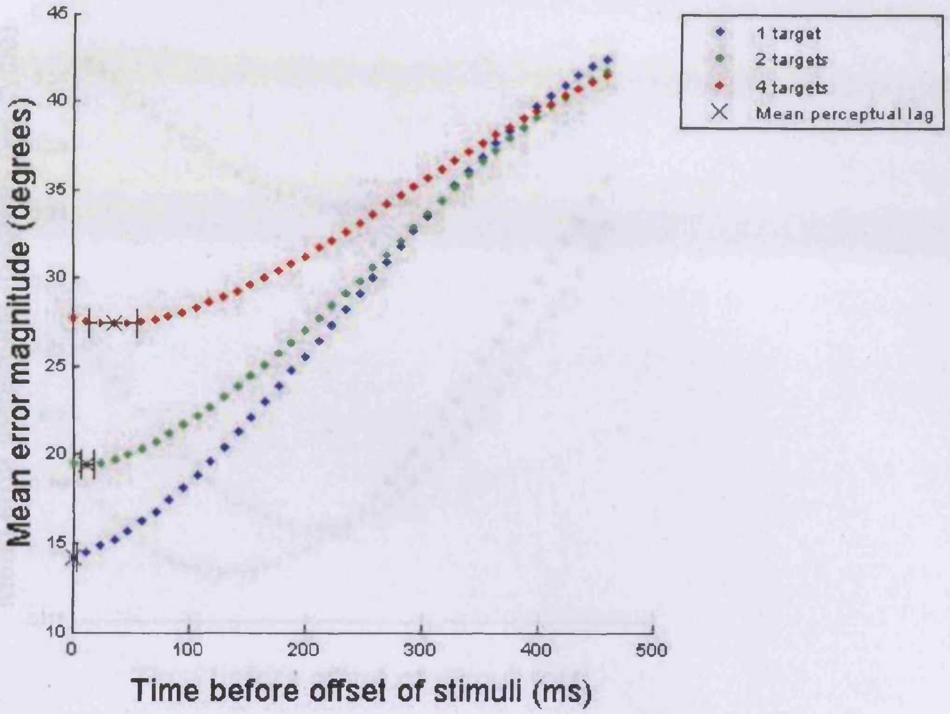

Figure 10: Lag analysis (Experiment 1: Orientation). Each individual dot represents the mean in the data of difference between the reported values and the value that the queried object had at each lag time before offset of the display. Mean perceptual lags are represented by minima on the curves of these differences. Note that the leftmost points represent the mean of the error magnitudes plotted in previous figures. Tracking more objects is associated both with greater mean error magnitudes and with greater perceptual lags. Error bars represent 1 SE.

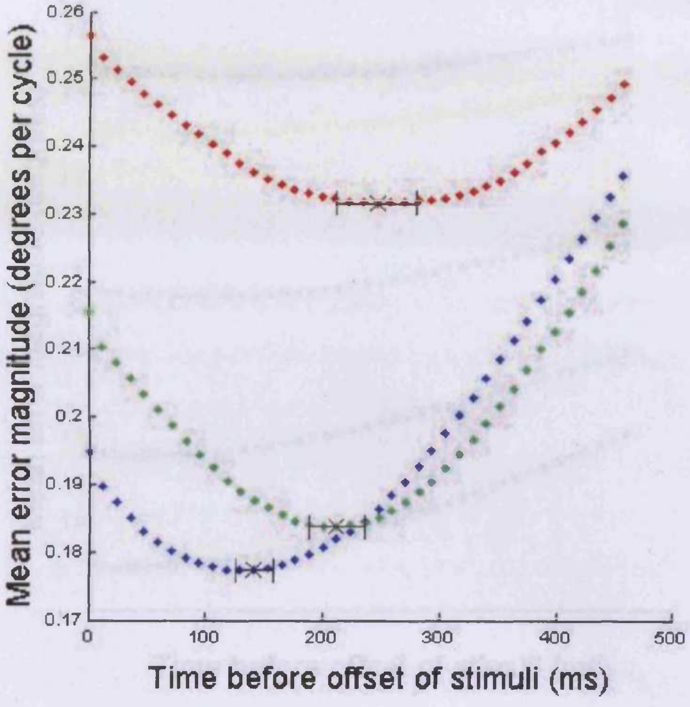

Figure 11: Lag analysis (Experiment 2: Spatial period). As in Experiment 1, tracking more objects is associated with greater mean error magnitudes and with greater perceptual lags. Tracking spatial periods is associated with large perceptual lags, much larger than those seen for tracking orientations in Figure 10. Error bars represent 1 SE.

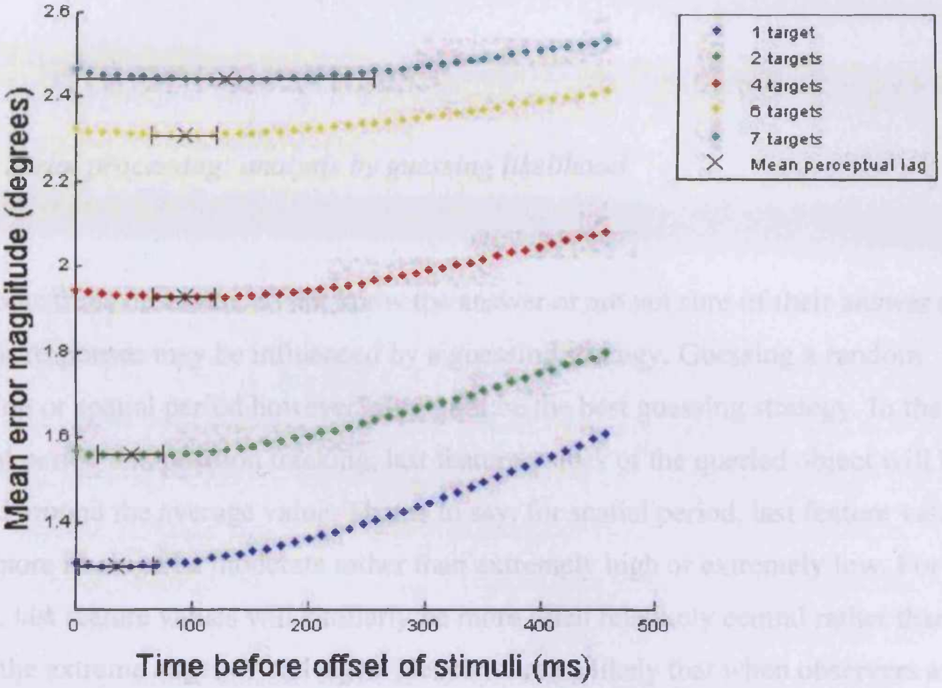

Figure 12: Lag analysis (Experiment 3: Position). As in Experiments 1 and 2, tracking positions shows a trend towards greater perceptual lags when tracking greater numbers of targets. Error bars represent 1 SE.

Plots of the data of individual observers are shown in Appendix 4 and reveal both differences in overall precision and differences in set size effects, although the general trend towards greater perceptual lags with greater numbers of objects attended can be seen.

Why would lag times increase with the number of objects tracked? There are three logical possibilities. The first possibility is that processing of all targets proceeds in parallel, but is progressively slowed the more objects are tracked. The second possibility is that processing proceeds in a serial manner, switching between target objects and also causing increasing lags in perception with more tracked objects. The third possibility is of course a combination of these two processes. The following analyses by guessing likelihood attempt to answer the question of what causes this pattern of set size effects on lags.

2.4.5 Serial processing: analysis by guessing likelihood

On some trials observers do not know the answer or are not sure of their answer and then their responses may be influenced by a guessing strategy. Guessing a random orientation or spatial period however would not be the best guessing strategy. In the case of spatial period and position tracking, last feature values of the queried object will be clustered around the average value. That is to say, for spatial period, last feature values will be more likely to be moderate rather than extremely high or extremely low. For position, last feature values will similarly be more often relatively central rather than lying at the extreme edges of ballcages. Hence it seems likely that when observers are less sure of their response, they will make a response on average relatively near to the average value. Furthermore, when they feel they do not know the answer, observers are likely to spend little time adjusting the sample and thus accept a value near the sample's starting value, which was the mean in the case of spatial period. Conversely, responses very far from the average value are less likely to be guesses. However, as orientation is a circular variable, on average no orientation was presented more often than any other, so there is no reason for guesses to have any particular distribution. Using this logic, apart from the case of orientation, we can conduct analyses directed at examining how responses with more contribution from guessing are different from those with less of this contribution from a guessing strategy.

If we select those trials from Experiment 2 where reported spatial period values are more extreme relative to the average value than the actual last states of the queried object, then these trials are likely to contain more trials where observers are confident in their responses, than for all trials on average. In other words, non-guessing responses should form a larger proportion of extreme responses than should guessing responses. The remaining trials, where reported values are less extreme than the actual last states of the queried Gabor, will contain more trials where observers are making some use of a guessing strategy.

Similarly for position tracking, trials were divided into those for which responses were spatially further from the centres of the imaginary cages than the actual answers on that trial (extreme responses), and those where responses were less extreme than the actual last positions of queried objects. Observers are likely to use a guessing strategy that is biased towards central values, which yields more guesses in the less extreme cases than in the more extreme cases.

Pure guesses will be clustered around the mean correct response, and hence will exhibit no particular lag at all on average. This makes sense, since all values of the queried object at different times are equally as likely to be similar to this average value. This logic is confirmed in section 2.4.7 in simulated data sets where all responses are modelled to occur at the average correct values. Guesses that are in part drawn from a guessing strategy, and in part from a relatively poor representation of the queried object can, however, display a lag corresponding to the time at which the representation was last updated.

2.4.6 Results of analysis by guessing likelihood: contribution from serial processing

The perceptual lag plots for these two groups of response types are shown in Figure 13 (spatial period tracking) and Figure 14 (position tracking). For those trials where responses were more extreme than the last value of the queried object in Experiment 2, there was no effect of attentional load on perceptual lags: perceptual lags were (mean±SE): 80±17 ms for tracking one Gabor, 120±19 ms for tracking two and 80±20 ms for tracking four. The percentages of trials falling into this category were 35.1%, 36.5% and 32.4%, respectively. With spatial period tracking, for the less extreme cases the observed lags were 150 ms for tracking one Gabor, 250 ms for tracking two and 350 ms for tracking four (errors bars represent standard errors: +- 42 ms for tracking one, +- 37 ms for tracking two, and +- 31 ms for tracking four). The percentages of the total number of trials falling into this 'less extreme' category were 59.4%, 57.9% and 62.1% for the three set sizes, respectively.

In Experiment 3 (position), for the more extreme cases, the lag pattern observed was 20±9 ms for tracking one, 40±26 ms tracking two, and 60±≈27 ms for tracking either four, six or seven objects. The frequencies of these more extreme cases were 39.9%, 34.6%, 31.2%, 32.1% and 32.8% respectively. For the less extreme (more contaminated by guessing) cases, the lag pattern observed was (mean±SE): 60±57 ms for tracking one, 80±39 ms for tracking two, 140±33 ms for tracking four, 150±133 ms for tracking six, and 200±129 ms for tracking seven objects (lags calculated from the following percentages of trials where responses were less extreme than the last state of the queried object: 60.1%, 65.4%, 68.8%, 67.9% and 67.2%, respectively).

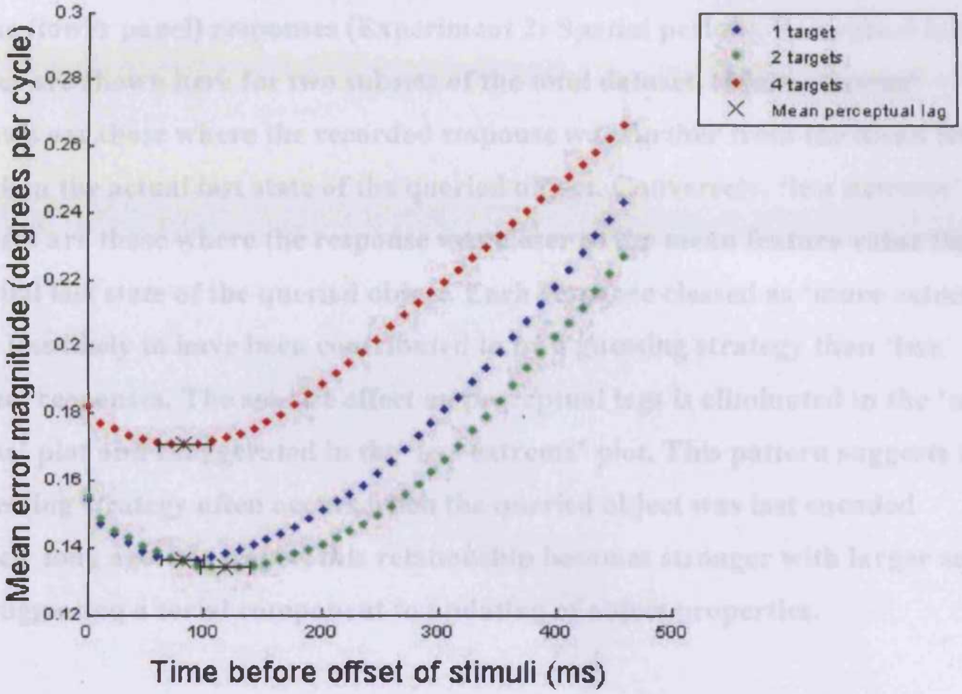
Mean error magnitude (degrees per cycle)
0.3
0.28
0.26
0.24
0.22
0.2
0.18
0.16
0.14
0.12
0
100
200
300
400
500
Time before offset of stimuli (ms)
1 target
2 targets
4 targets
Mean perceptual lag

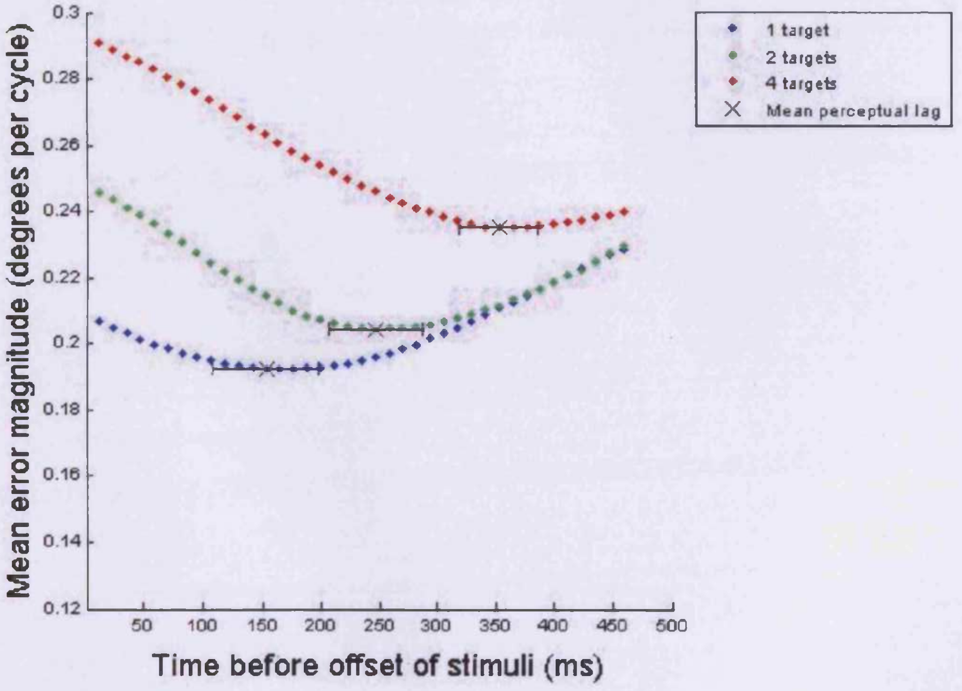
Mean error magnitude (degrees per cycle)
0.3
0.28
0.26
0.24
0.22
0.2
0.18
0.16
0.14
0.12
50
100
150
200
250
300
350
400
450
500
Time before offset of stimuli (ms)
1 target
2 targets
4 targets
Mean perceptual lag

Figure 13: Mean absolute errors for more extreme (upper panel) and less extreme (lower panel) responses (Experiment 2: Spatial period). Perceptual lag analyses are shown here for two subsets of the total dataset. 'More extreme' responses are those where the recorded response was further from the mean feature value than the actual last state of the queried object. Conversely, 'less extreme' responses are those where the response was closer to the mean feature value than the actual last state of the queried object. Each response classed as 'more extreme' will be less likely to have been contributed to by a guessing strategy than 'less extreme' responses. The set-size effect on perceptual lags is eliminated in the 'more extreme' plot and exaggerated in the 'less extreme' plot. This pattern suggests that the guessing strategy often occurs when the queried object was last encoded relatively long ago. Moreover, this relationship becomes stronger with larger set-sizes, suggesting a serial component to updating of object properties.

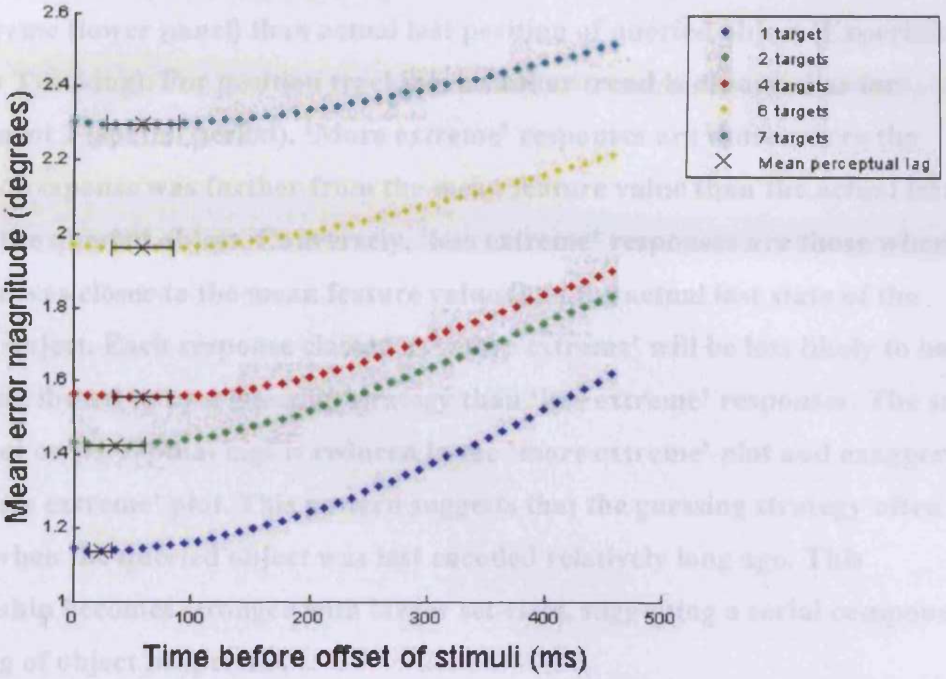
2.6
2.4
2.2
2
1.8
1.6
1.4
1.2
1
Mean error magnitude (degrees)
0
100
200
300
400
500
Time before offset of stimuli (ms)
1 target
2 targets
4 targets
6 targets
7 targets
Mean perceptual lag

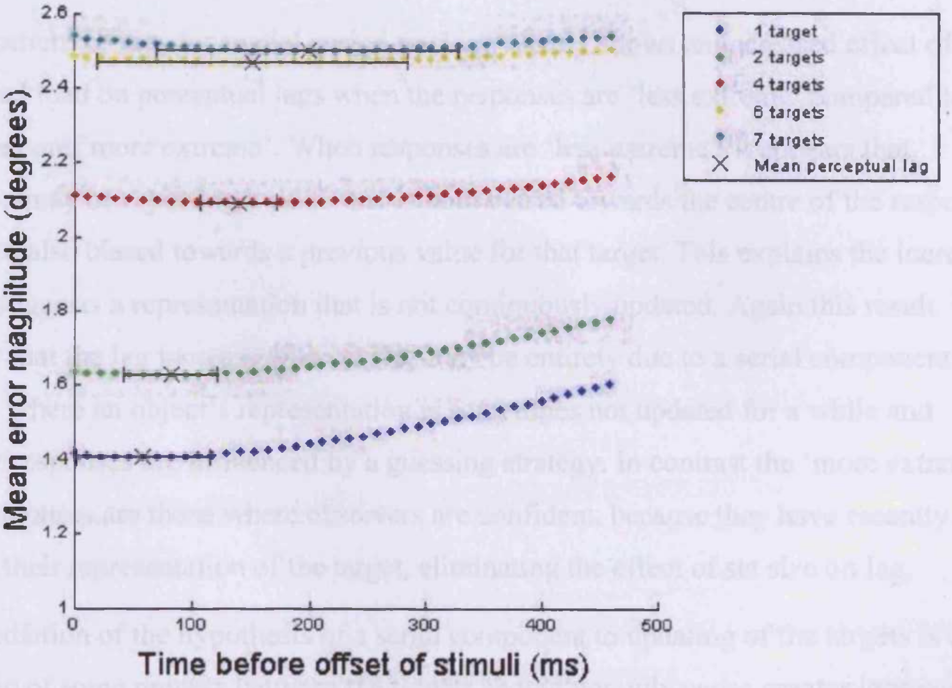
2.6
2.4
2.2
2
1.8
1.6
1.4
1.2
1
Mean error magnitude (degrees)
0
100
200
300
400
500
Time before offset of stimuli (ms)
1 target
2 targets
4 targets
6 targets
7 targets
Mean perceptual lag

Figure 14: Perceptual lag analysis for more extreme (upper panel) responses and less extreme (lower panel) than actual last position of queried object (Experiment 3: Position Tracking). For position tracking, a similar trend is observed as for Experiment 2 (spatial period). 'More extreme' responses are those where the recorded response was further from the mean feature value than the actual last state of the queried object. Conversely, 'less extreme' responses are those where the response was closer to the mean feature value than the actual last state of the queried object. Each response classed as 'more extreme' will be less likely to have been contributed to by a guessing strategy than 'less extreme' responses. The set-size effect on perceptual lags is reduced in the 'more extreme' plot and exaggerated in the 'less extreme' plot. This pattern suggests that the guessing strategy often occurs when the queried object was last encoded relatively long ago. This relationship becomes stronger with larger set-sizes, suggesting a serial component to updating of object properties.

The pattern of lags for spatial period tracking clearly shows an increased effect of attentional load on perceptual lags when the responses are 'less extreme' compared to when they are 'more extreme'. When responses are 'less extreme', it appears that observers may be reporting a value that is both biased towards the centre of the response range and also biased towards a previous value for that target. This explains the increased lag and suggests a representation that is not continuously updated. Again this result suggests that the lag increase with set size may be entirely due to a serial component to the task, where an object's representation is sometimes not updated for a while and observer responses are influenced by a guessing strategy. In contrast the 'more extreme' set of responses are those where observers are confident, because they have recently updated their representation of the target, eliminating the effect of set size on lag.

A prediction of the hypothesis of a serial component to updating of the targets is that switching of some process between the targets should not only cause greater lags with greater numbers of targets, but also greater variability in lags observed for greater numbers of targets. In the case of two targets, in the extreme case where only one object is updated at a time, then depending on which was queried, we would expect either a lag near zero or a lag as long as the duration for processing both objects. As each target is added to the to-be-attended set, there should be an additional possible amount of time since an item was last updated, increasing the variability of lags and the difference between the lags for 'less extreme' and 'more extreme' responses. For the position tracking data there appears to be the same trend for a greater effect of attentional load on perceptual lags for less extreme cases, and correspondingly less effect for more extreme responses, although this pattern is less pronounced than was observed for spatial period tracking. The differing size of the trend across experiments might result from differing extent to which a serial component was present, or perhaps reflect different capacities for the parallel process involved.

The trends described in the previous paragraph are largely confirmed by correlations performed between observed lag times and set sizes. Significant correlations were observed for both spatial period (E2) and position (E3) when analysing the lag patterns for the 'less extreme' responses (spatial period, E2: correlation coefficient = 0.714, N = 18, p (1 tailed) < 0.01; position, E3: correlation coefficient = 0.414. N = 35, p (1 tailed) < 0.01). Conversely, correlation analyses were non-significant for both sets of 'more extreme' data (spatial period, E2: correlation coefficient = -0.186, N = 18, p (1 tailed) = 0.230; position, E3: correlation coefficient = 0.269, N = 35, p (1 tailed) = 0.059). In analyses of variance, the interaction between the effects of set size and whether the responses used were 'less extreme' or 'more extreme', was indeed significant for spatial period (E2: F (2, 30) = 7.778, $p < 0.01$) indicating a strong effect of set size on perceptual lags only for 'less extreme' responses. For position monitoring, this trend may be present but it was non-significant (E3: F (4, 60) = 1.365, p = 0.257).

The pattern of lags that we observe when broken down by guessing likelihood is predicted straightforwardly from a serial processing model. A parallel processing model can only account for this pattern with two strong assumptions about the architecture of the processing system: it is not, therefore, incompatible with the data, but the serial model is the more parsimonious explanation. A parallel model would need to account for the elimination of the set size effect for the responses that are less likely to contain a high proportion of guesses. This requires the assumption that the parallel system architecture permits a ceiling level of performance where the set size effect on lags is eliminated. An alternative conceptualisation of the parallel model could predict that no elimination of the set size effect could ever be achieved, if additions to the set size necessarily produce longer lags.

A further assumption is required to explain the exaggerated set size effect for the responses that are more likely to contain a high proportion of guesses. The assumption required is that the parallel processing system suffers more in terms of processing time with each addition to set size, the greater the lag time. An alternative conceptualisation of the parallel model could conceivably predict that additions to the set size should add a constant amount to the lag time, regardless of the absolute values of the lag. Since we do not know whether either of these assumptions are valid, the serial model appears to be the more parsimonious of the two explanations for the patterns observed in lags when broken down by guessing likelihood.

2.4.7 Eliminating explanations based on statistics of trajectories

To confirm that these patterns of lags could not be artefacts stemming from statistical patterns in the trajectories of objects, we also ran these analyses using simulated dummy sets of data (for more details see Appendix 9). First, for Experiment 2 we analysed lag patterns for simulated datasets where all 'responses' were simulated to occur at the mean feature value. In addition lag patterns were calculated for responses simulated to occur at artificially extreme values. For Experiment 3, the analyses included a data set where all 'responses' were simulated to occur at the mean recorded response location for that cage, as this is likely to be close to the most frequent location used as a guess. In addition, we used a simulated data set where all 'responses' were recorded as being near the most peripheral edge of cages, as this is a statistically unlikely location of a guess response. No differences in lags were observed in these simulations. Indeed the appearance of perceptual lags cannot arise from guessing, since guessing only serves to increase the mean error magnitude, and not the time at which the feature value of the queried object was most similar to the reported value.

2.4.8 Summary of guessing likelihood analysis: implications for serial processing

The outcomes of these analyses are consistent with observers switching their processing resources between targets for tracking, such that lag times are at a constant and small value when the queried object is the last-processed object or objects. When the queried object is not likely to have been the last-processed object, lags are increasingly longer the more objects must be attended, implicating a serial process where the lag reflects the last time on average that the object was updated. Of course, this updating could reflect a serial process at either the attentional or memory level.

Another prediction of serial processing arises if observers are asked to report multiple targets on individual trials, and this is investigated in Experiments 4 and 5.

CHAPTER 3. EXPERIMENTS 4 & 5: DISTINGUISHING BETWEEN PARALLEL AND SERIAL PROCESSING

3.1 INTRODUCTION

3.1.1 Overview of Experiments 4 and 5

Experiments 4 and 5 were designed to allow us to compare the relationship between performance in reporting each of two target Gabors, out of a total of either two or four tracked Gabors. In these experiments, observers first reported one target in exactly the same way as in Experiments 1 and 2 but then were immediately cued again to report a second target. We inferred that serial and parallel processing should cause different signature patterns in the correlations between accuracy in reporting a first-queried and a second-queried target. These different patterns could then allow us to differentiate between likely serial or parallel processing, or a combination of the two.

3.1.2 Rationale for Experiments 4 and 5: distinguishing between serial and parallel processing

Because in this experiment observers must report the features of two targets, there are two new sources of imprecision expected in performance. First, accuracy during reporting the first Gabor may suffer because observers have the concurrent demand of holding in memory the value of the other monitored object(s). Second, they may be more likely to have a degraded representation in memory of the second reported object because of the additional time before reporting can occur. However, these sources of imprecision, though increasing error sizes on both first and second reports, cannot affect correlation patterns between pairs of errors.

If observers are rapidly switching processing between monitored objects in a serial manner, then we would expect a negative correlation between accuracy on reporting the first and second targets, as processing resources given to one will necessarily be associated with it being withdrawn from the other target or targets.

3.1.3 Worked example

If targets are serially processed, on a trial with two target objects, attention would first be directed in full to one object, then after a brief period it would be withdrawn and redirected to the other. Consider an example trial where there are two targets, one on the left and the other on the right-hand side of the screen. Since observers cannot know in advance when or which object will be first queried, there is no reason for one object to be favoured over the other. At the moment the objects disappear, attention will have last been directed to the left-hand object on 50% of trials and to the right-hand object on the other 50% of trials. In both these scenarios, the first queried object will be on the left-hand side in 50% of trials, and on the right in the other 50% of trials. Thus, there are four possible (2x2) combinations of which object was last attended and which was first queried: let us consider each in turn.

In scenario 1, the right-hand target was last attended, and was also queried first. In this scenario, we would expect relatively good performance for reporting the right-hand target, and relatively poor performance for reporting the left-hand target. In scenario 2, the right-hand target was last attended, but was queried second. In this scenario, the right-hand target will still benefit from the attentional resource it received, to an extent inversely proportional to the resources received by the left-hand target. Scenarios 3 and 4 are, of course, the equivalents of scenarios 1 and 2, except that the left-hand target is the one that received the benefit of full attentional resources at the moment the objects disappeared. Thus, each pair of reports on every trial should reflect the negative relationship between processing being directed to one object and it being directed to the other, regardless of whether the last attended target was reported first or second.

If observers are sharing a processing resource over multiple targets in a parallel manner simultaneously, then one would expect a positive correlation between accuracy of reporting the first and second targets. The source of variability in performance that leads to the correlation might be fluctuations from trial to trial in the amount of attentional resource available, or fluctuations in general arousal. On any particular trial, one would expect this attentional resource to be directed to all targets to a similar extent, since observers have no idea which of the targets will be queried first and second.

If there is no correlation, then this is consistent with contributions from both switching and sharing of processing between targets for tracking. This could result, for instance, from relatively more switching on trials where there is less attentional resource available (perhaps due to distraction, less concentration or less general arousal).

Our double report correlations are similar to the methods used by Sperling and Melchner (1978) in their 'attention operating characteristic' (AOC). Their task was a visual search amongst two concentrically arranged arrays of stimuli. Observers were sometimes instructed to devote 90% of their attention to the outer array of stimuli, and sometimes, conversely, to devote 90% of their attention to the inner array. Sperling and Melchner then examined the trade-off between detection probabilities for each array. As a control measure, they also measured performance when only targets in one or the other array were to be reported. They inferred that pure independence of attention to one or the other array should result in no decrement in detection probabilities when observers attended to the two arrays, compared to the control levels of performance. Any drop in performance for the two arrays when attended simultaneously would represent the degree to which the two tasks interfere with one another. From trial to trial, if observers vary the amount of their processing resources directed to one of the arrays relative to the other, they predict a strict interdependence between performance for the two arrays i.e. a negative correlation between performance for the two arrays. Our analysis shares this assumption that if two tasks draw on the same resource, attention directed to one should directly cause it to be directed away from the other, causing a trade-off in performance.

3.2 METHOD FOR EXPERIMENTS 4 AND 5

The procedures for Experiments 4 and 5 were identical to Experiments 1 and 2, except that observers reported the last feature values of two separate targets for tracking on each trial. In these experiments, the number of Gabors to be tracked on any trial was either two or four. Observers reported one feature value exactly as in Experiments 1 and 2, they were then prompted in the same way to report either the other of the pair (in the tracking two Gabors condition) or one of the remaining three (in the tracking four Gabors condition). The post-cue for the second report appeared immediately when the next trial would have been presented in Experiments 1 and 2 after the observer made a response.

3.2.1 Observers

Observers were five postgraduate students at Cardiff University and one undergraduate student. Four were naïve as to the purpose of the experiment, and four were non-experts in participating in visual psychophysical experiments. All had normal or corrected-to-normal vision.

3.2.2 Design

For both experiments, observers completed three blocks of 70 trials. Within each block there were 35 trials for tracking two objects and 35 trials for tracking four objects, yielding a total of 105 trials for each condition for each observer. On each trial, the number of targets for tracking was selected randomly until for each block, each of the two conditions (tracking two and tracking four) had been run 35 times.

3.3 RESULTS AND DISCUSSION FOR EXPERIMENTS 4 AND 5

Scatterplots relating error magnitudes on the first and second reports are shown in Figure 15 (Experiment 4: orientation) and Figure 16 (Experiment 5: spatial period). When reporting two orientations (Figure 15), there is a small but significant positive correlation between errors recorded during the first and second reports (correlation coefficient = 0.116, N = 1260, p (2-tailed) < 0.01). This is consistent with parallel processing of the attended Gabors. Variation in overall accuracy on the two reports could be contributed to by factors such as levels of overall alertness or concentration on that trial. The inter-error correlation is also small but significant for the individual conditions of tracking two Gabors (correlation coefficient = 0.108, N = 630, p (2-tailed) < 0.01) and for those tracking four Gabors (correlation coefficient = 0.099, N = 630, p (2-tailed) < 0.05). These results for tracking orientations are not consistent with a low-capacity serial process, which would be expected to yield a large negative correlation. These small correlations could result from an entirely parallel or more likely a limited-capacity parallel process with a serial model, contributing to both positive and negative correlations that might mostly cancel out, as may have occurred here.

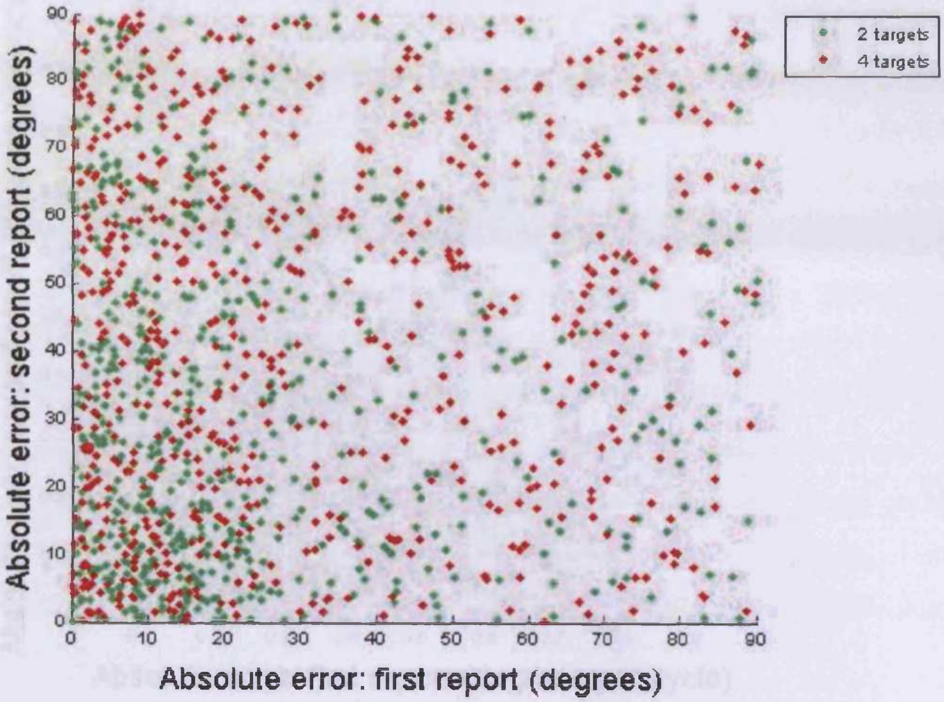

Figure 15: Errors in reporting last orientation of first queried and second queried Gabors. Errors in reporting the orientation for the first and second queried targets are correlated, although they are well spread over the total space of possible errors. This is true of the total dataset, and also for each of the two set size conditions separately. Not surprisingly, first reported targets are associated with smaller error magnitudes than second reported targets, presumably as the effect of memory decay of the second target is stronger than the effect of the requirement to perform a concurrent task (to hold in memory a second response) whilst making the first response. A positive correlation is consistent with parallel processing of targets, since accuracy on both reports is affected equally by the same availability of the same attentional resource on that trial. Fluctuations in the level of attentional resource from trial-to-trial could arise, for instance, through fluctuations in concentration or in general arousal.

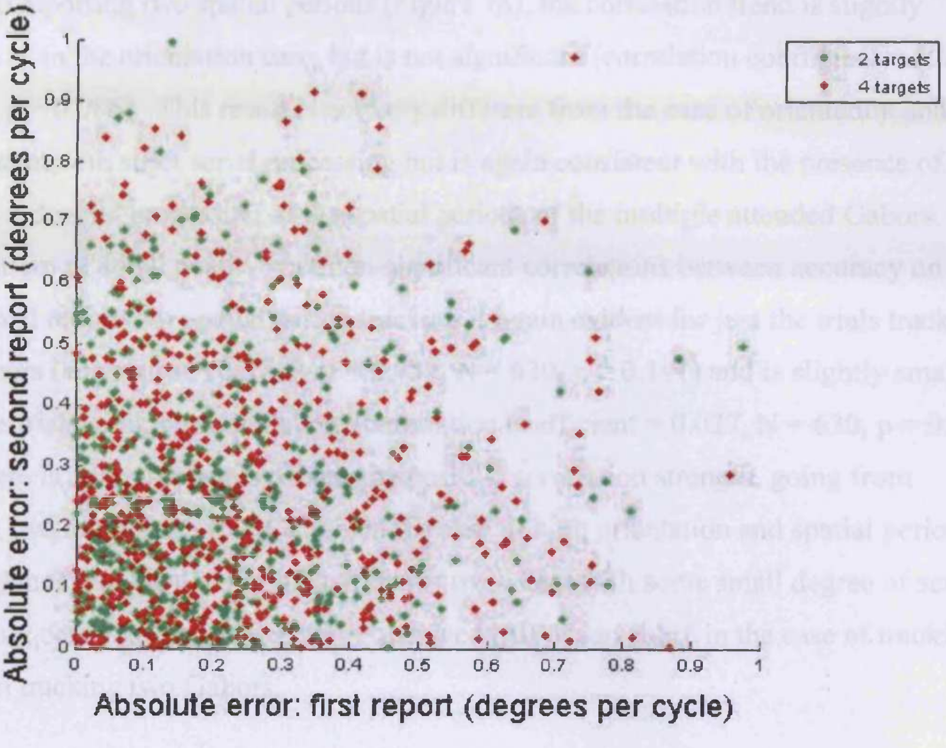

Figure 16: Errors in reporting last spatial period of first queried and second queried Gabors. For spatial period tracking, no correlation is observed between accuracy of reporting the first queried target and accuracy during the report of the second queried target. This is true of the total dataset, and also for each of the two set size conditions separately. Although care should of course be taken in drawing conclusions from a null result, no correlation is logically consistent with the presence of both serial and parallel processing elements. A positive correlation would result from a purely parallel model where reports of both targets are affected equally by the same level of availability of attentional resources. A negative correlation would result from a purely serial model where attention to one object necessarily entails withdrawal of attention from the other. A hybrid model with both serial and parallel components could produce no overall correlation, as is evident here.

When reporting two spatial periods (Figure 16), the correlation trend is slightly positive as in the orientation case, but is not significant (correlation coefficient = 0.048, N = 1260, p = 0.086). This result is not very different from the case of orientation and is inconsistent with strict serial processing but is again consistent with the presence of both parallel and serial processing of the spatial periods of the multiple attended Gabors. The same pattern of small positive but non-significant correlations between accuracy on first and second reports for spatial period tracking is again evident for just the trials tracking two Gabors (correlation coefficient = 0.059, N = 630, p = 0.141) and is slightly smaller for those trials tracking four Gabors (correlation coefficient = 0.027, N = 630, p = 0.493). Thus there is a trend towards decreasing positive correlation strength, going from tracking two to tracking four Gabors in the case of both orientation and spatial period. Although not statistically significant, this is consistent with some small degree of serial processing contributing to the pattern observed that occurs more in the case of tracking four than tracking two Gabors.

3.3.1 Summary of results from Experiments 4 & 5: Distinguishing between serial and parallel processing

In Experiments 4 and 5, observers continuously monitored the changing orientations and spatial periods, respectively, of multiple targets. The methods used were identical to those in Experiments 1 and 2, except that observers reported two out of a total of either two or four target objects. These experiments were designed to distinguish between serial and parallel processing of target objects, since the two processes were predicted to cause different signature patterns in the correlations between error magnitudes on the first and second reports. The effects of the additional task demands of making two reports on every trial were considered to increase errors on both the first and second reports, but could not affect correlations between reports.

A positive correlation would result from a purely parallel model where reports of both targets are affected equally by the same level of availability of attentional resources. This positive correlation was observed in the data for orientation monitoring (Experiment 4).

A negative correlation would result from a purely serial model where attention to one object necessarily entails withdrawal of attention from the other. A hybrid model with both serial and parallel components could produce no overall correlation, as was evident for Experiment 5 (spatial period tracking). A null correlation does of course mean that drawing a conclusion about the data is risky, but we can say that the data are logically consistent with a hybrid process containing serial and parallel components. This is also consistent with the large set size effects on lag times observed previously in Experiment 2 for spatial period tracking, and also with the lag analysis broken down by guessing likelihoods, which were consistent with the presence of a serial component to processing, especially in the case of spatial period tracking.

To investigate further why spatial period monitoring appears to exhibit such large lags, such a large set size effect on lags, and a potentially serial component to processing, we considered in Chapter 4 whether there were artefacts in the nature and range of feature values used in spatial period tracking that were different from those for position and orientation monitoring. In four control experiments in Chapter 4, we manipulate the feature values used for orientation monitoring to attempt to simulate those previously used for spatial period. None of these manipulations will lead to the patterns of lags observed for spatial period monitoring, leading us to conclude that it was the feature of spatial period itself, which led to the lag patterns that were so different from orientation and position monitoring. In Chapter 5 we discuss potential reasons for these differences between monitoring different features.

CHAPTER 4. EXPERIMENTS 6 – 9: CONTROL EXPERIMENTS FOR POSSIBLE ARTEFACTS IN THE RANGE OF FEATURE VALUES

4.1 INTRODUCTION

4.1.1 Rationale for Experiments 6 -9

Experiments 6 – 9 were designed to investigate the reason for the difference in lag results between orientation and spatial period monitoring. To anticipate, we will test four possible reasons for the difference between the lag patterns observed for spatial periods and the other two features tested; orientation and position. None of the four control experiments presented in Chapter 4 will account for the difference, and hence we will conclude that the difference in lag patterns observed for spatial period is due to an inherent difference in the way that spatial periods were monitored, and not simply due to characteristics of the spatial period feature range used. In Chapter 5, possible reasons for the difference between spatial period and the other two features used (orientation and position) are discussed.

There are several differences between the stimuli in Experiments 1 and 2 other than the mere fact of monitoring different features. Firstly, orientation is a circular variable and spatial period is a linear variable. As the difference between any two orientation values increases, they will eventually start to become more similar, until the difference between the values is 360 degrees. This is not the case for spatial period values which become more different the larger the separation between them. One result of this difference, is that in Experiment 2, the range of possible spatial period values had a maximum and a minimum, in contrast to the orientation values which had no limits i.e. they could hypothetically spin indefinitely. Secondly, the discriminability of the changes in feature values may have been different in the two experiments. 70% correct thresholds for discriminating the difference between the spatial periods of two sine wave gratings are in the order of 0.1 dpc, when the spatial period values of the two gratings are around 1 dpc and for presentation both at the fovea and at 10 degrees eccentricity (stimulus values similar to the Gabor patches used in Experiment 2 here), (Richter and Yager, 1984). This threshold equates to approximately one tenth of the total range (1.1 dpc) of spatial periods used in Experiment 2. Thresholds for discriminating the orientations of high contrast gratings are in the order of 0.15 - 0.5 degrees (c.f. Regan, 2000) and are less than 0.05 degrees for lines approximately 3 degrees long and presented at approximately 7 degrees eccentricity (stimulus values similar to the Gabor patches used in Experiment 1 here), (Mäkelä, Whitaker & Rovamo, 1993). Even at the least discriminable end of this range, these thresholds represent approximately less than one percent of the total range of orientation values presented in Experiment 1 (180 degrees). The changes in spatial period presented in Experiment 2 were thus much less discriminable than the changes in orientation presented in Experiment 1. A result of this difference is that the perceived rate of change of values is also likely to have been different.

4.1.2 Overview of Experiments 6 - 9

It seems possible that any of these differences could have contributed to the difference in lag results found between Experiments 1 and 2. Experiments 6 – 9 were designed to rule out these possibilities by examining the effects of various feature parameters on perceptual lags for orientation monitoring. In Experiment 6, we investigate the effect of range of feature values. In Experiment 7 we investigate the effect of discriminability of feature changes and in Experiments 8 and 9 we examine the effect of the rate of feature change. The lag analysis itself should not be affected by the rate of feature change. For instance, if perceptual lags were constant, a faster rate of feature change would reveal the same perceptual lag as a slower rate of feature change, since with a given perceptual lag of x ms, observers will still report the feature value from approximately x ms in the past, regardless of how rapidly the object has been changing in the intervening period. Hence, any difference in perceptual lags would have to be attributed to other factors such as a change in object updating by the visual system. If, for some reason, either the range or discriminability of feature values or the rate of feature change affect this object updating, then these experiments should reveal this fact and hence start to explain the difference in lag results observed between Experiments 1 and 2.

4.2 METHOD FOR EXPERIMENTS 6 – 9

Experiments 6 – 9 were identical to Experiment 1 except for the following details. In Experiment 6, as in Experiment 1, initial starting orientations for each Gabor were each randomly chosen from between 0 and 359 degrees. In Experiment 6, however, if the orientation of any Gabor reached either 0 or 360 degrees, it simply 'bounced' off this mimimum or maximum value by changing the direction of its spin and thereby producing sharp reversals in orientation change. This made the pattern of orientation changes more comparable to that of the spatial frequency changes in previous experiments.

Experiment 7 was exactly the same as Experiment 6, except that the range of potential orientation values was restricted to between 0 and 45 degrees, corresponding to between vertical and tilted 45 degrees to the left. This made the range more comparable to the spatial period experiments, which perceptually spanned over the order of 10 multiples of discrimination threshold. Whilst the range of orientations in Experiment 7 was still more discriminable than that used for spatial periods in Experiment 2, pilot tests showed that ranges smaller than 45 degrees resulted in very high levels of noise in the data. Experiment 7 was thus conducted to investigate the contribution of lowered discriminability of feature changes over and above the addition of feature boundaries introduced in Experiment 6.

Experiments 8 and 9 were identical to Experiment 1 except for the maximum angular velocities allowed and the possible angular acceleration values. In Experiment 8, Gabors rotated more slowly and had smaller absolute accelerations than in Experiment 1. Conversely, in Experiment 9, Gabors rotated faster on average, and were given greater absolute accelerations than in Experiment 1. In Experiment 8, the maximum absolute angular velocity was only 0.09 degrees per ms, and in Experiment 9 the maximum was

0.43 degrees per ms. In Experiment 8, the pair of possible angular accelerations were -3.61×10^{-5} deg per ms^2 or 3.61×10^{-5} deg per ms^2 degrees per $frame^2$. If the angular velocity was below 0.017 degrees per ms, then the absolute value of the acceleration was increased from 3.61×10^{-5} deg per ms^2 to 1.45×10^{-4} deg per ms^2. In Experiment 9, the pair of possible angular accelerations were -1.08×10^{-4} deg per ms^2 or 1.08×10^{-4} deg per ms^2 degrees per $frame^2$. If the angular velocity was below 0.017 degrees per ms, then the absolute value of the acceleration was increased from 1.08×10^{-4} deg per ms^2 to 2.89×10^{-4} deg per ms^2.

4.2.1 Observers

Observers in Experiment 6 were five postgraduate students at Cardiff University and one undergraduate student. Four were naïve as to the purpose of the experiment, and two were non-experts in participating in visual psychophysical experiments.

Observers in Experiment 7 were six postgraduate students at Cardiff University. Four were naïve as to the purpose of the experiment, and two were non-experts in participating in visual psychophysical experiments.

Observers in Experiment 8 were three postgraduate students at Cardiff University and three undergraduate students. Four were naïve as to the purpose of the experiment, and four were non-experts in participating in visual psychophysical experiments.

Observers in Experiment 9 were five postgraduate students at Cardiff University and one undergraduate student. Three were naïve as to the purpose of the experiment, and three were non-experts in participating in visual psychophysical experiments.

All observers in Experiments 6 -9 had normal or corrected-to-normal vision.

4.3 RESULTS AND DISCUSSION OF EXPERIMENTS 6 – 9

4.3.1 Error analyses

Two analyses were performed on the data from Experiments 6 – 9: analyses of error magnitudes and lag patterns. We present the error analyses first. Error variances and mean magnitudes for Experiments 6 – 9 are shown in Table 2. Error variances were smallest for Experiment 7, which is not surprising, as the total range of possible orientations were severely restricted. Indeed in Experiment 7, for each of the three attentional load conditions (tracking one, tracking two and tracking four objects), error variances were significantly different from those observed in Experiment 1.

All four experiments show significantly different error variances (set size effects) between each of the three attentional load conditions.

Error variances were significantly greater for tracking one object in Experiment 6 than in Experiment 1. In Experiment 8, error variances were significantly greater than those in Experiment 1 for tracking one and tracking two objects. In Experiment 9, error variances were significantly greater than in Experiment 1 for tracking one object. It seems possible that the sharp reversals in orientation change in Experiment 6 contributed to difficulty in accurately reporting orientations, perhaps through making extrapolatory mechanisms less effective. Performance in terms of mean error magnitude was worse by 4.3 degrees and 2.5 degrees respectively in Experiments 8 and 9, with slower and faster speeds of orientation change, respectively. It seems possible that the speed of feature change relates to performance in a relatively complex way. It may be that faster speeds are more difficult to continuously monitor, but also benefit from greater exogenous attentional capture from luminance-defined motion signals.

Experiment 6 (orientation bounces off 0 and 360)	Error variance (deg^2)	Mean error magnitude (deg)	Significantly different from Experiment 1 equivalent?
Tracking one	662	17.9	Levene statistic = 15.549 d.f. = 1, 1258 $p < 0.01$
Tracking two	906	20.7	Levene statistic = 1.589 d.f. = 1, 1258 $p = 0.208$
Tracking four	1382	27.7	Levene statistic = 0.232 d.f. = 1, 1258 $p = 0.630$
Difference between one and two	Levene statistic = 6.057 d.f. = 1, 1258 $p < 0.05$		
Difference between one and four	Levene statistic = 64.038 d.f. = 1, 1258 $p < 0.01$		
Difference between two and four	Levene statistic = 28.694 d.f. = 1, 1258 $p < 0.01$		
Experiment 7 (orientation bounded by 0 and 45 deg)	Error variance (deg^2)	Mean error magnitude (deg)	Significantly different from Experiment 1 equivalent?
Tracking one	84	6.96	Levene statistic =121.01 d.f. = 1, 1258 $p < 0.01$
Tracking two	123	8.83	Levene statistic = 145.244 d.f. = 1, 1258 $p < 0.01$
Tracking four	187	10.6	Levene statistic = 248.631 d.f. = 1, 1258

			$p < 0.01$
Difference between one and two	**Levene statistic = 20.399 d.f. = 1, 1258 $p < 0.01$**		
Difference between one and four	**Levene statistic = 60.369 d.f. = 1, 1258 $p < 0.01$**		
Difference between two and four	**Levene statistic = 13.748 d.f. = 1, 1258 $p < 0.01$**		
Experiment 8 (slow rotation)	**Error variance (deg^2)**	**Mean error magnitude (deg)**	**Significantly different from Experiment 1 equivalent?**
Tracking one	710	18.4	**Levene statistic = 19.506 d.f. = 1, 1258 $p < 0.01$**
Tracking two	947	21.7	**Levene statistic = 4.443 d.f. = 1, 1258 $p < 0.05$**
Tracking four	1174	24.7	**Levene statistic = 3.002 d.f. = 1, 1258 $p = 0.083$**
Difference between one and two	**Levene statistic = 7.984 d.f. = 1, 1258 $p < 0.01$**		
Difference between one and four	**Levene statistic = 26.188 d.f. = 1, 1258 $p < 0.01$**		
Difference between two and four	**Levene statistic = 5.308 d.f. = 1, 1258 $p < 0.05$**		
Experiment 9	**Error**	**Mean**	**Significantly different from**

(fast rotation)	**variance (deg^2)**	**error magnitude (deg)**	**Experiment 1 equivalent?**
Tracking one	616	16.6	**Levene statistic = 7.117** **d.f. = 1, 1258** **p < 0.01**
Tracking two	907	20.6	**Levene statistic = 1.408** **d.f. = 1, 1258** **p = 0.236**
Tracking four	1126	24.5	**Levene statistic = 3.542** **d.f. = 1, 1258** **p = 0.060**
Difference between one and two	**Levene statistic = 11.970** **d.f. = 1, 1258** **p < 0.01**		
Difference between one and four	**Levene statistic = 44.948** **d.f. = 1, 1258** **p < 0.01**		
Difference between two and four	**Levene statistic = 9.520** **d.f. = 1, 1258** **p < 0.01**		

Table 2: Comparisons of error distributions in Experiments 6 - 9 with Experiment 1

Error distributions for the three different set sizes, and for each of Experiments 6 – 9 are shown in Figures 17 - 20 below.

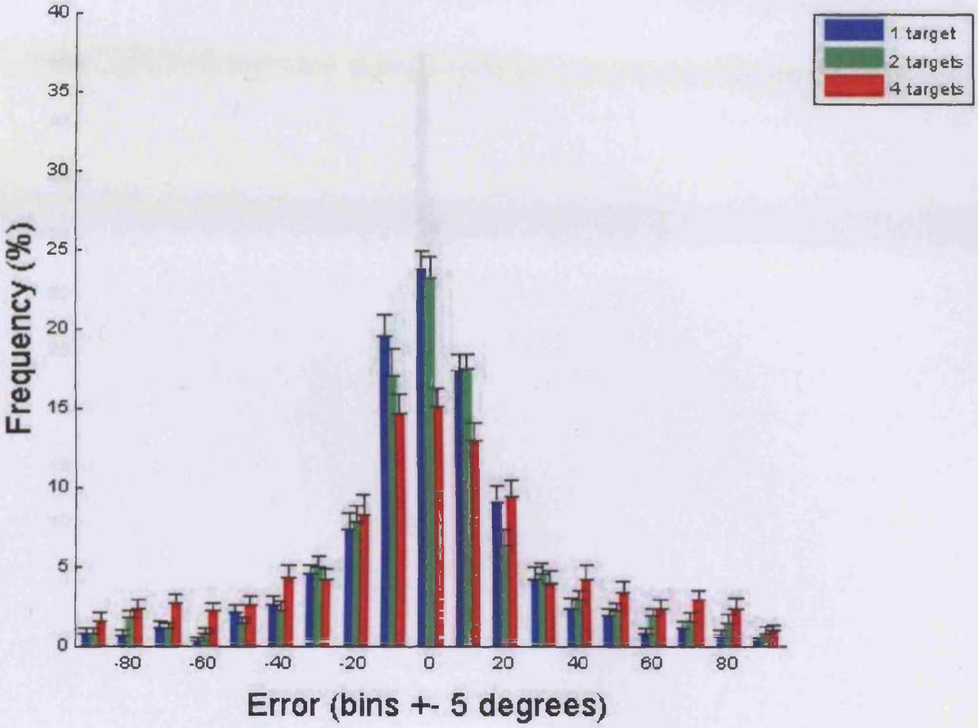

Figure 17: Errors in reported orientation for Experiment 6. This experiment included sharp reversals when orientations reach 0 or 360 degrees. Error bars represent 1 SE. N = 6. Although error magnitudes are different from those seen in Experiment 1, the same pattern is observed as was seen in Experiment 1, namely a gradual increase in spread of errors with the number of objects tracked.

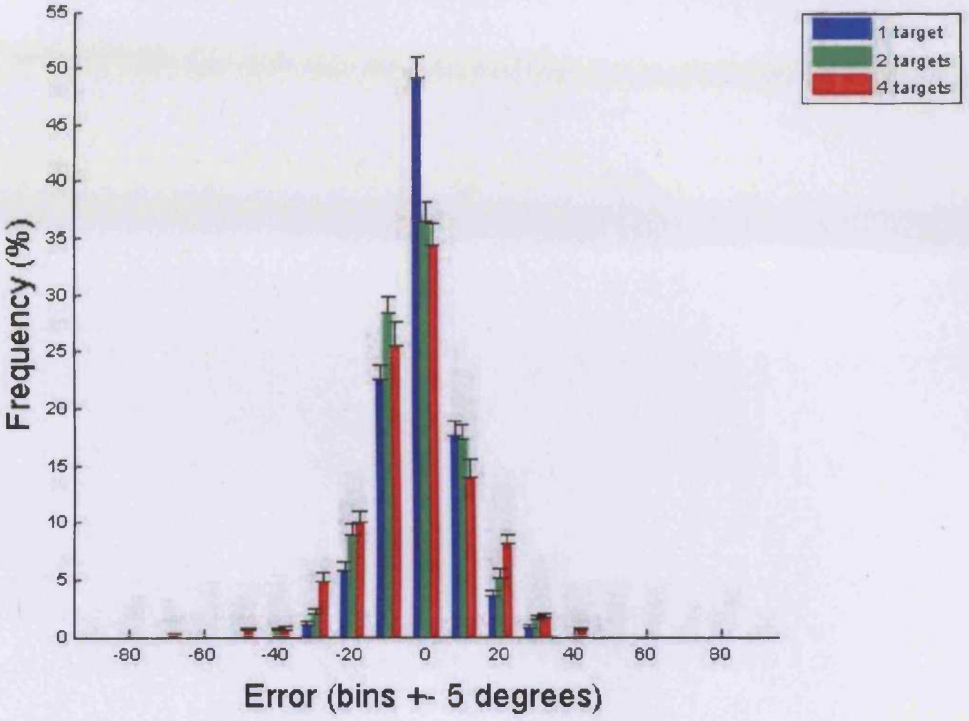

Figure 18: Errors in reported orientation for Experiment 7, for which the orientation range was restricted to between 0 and 45 degrees. Error bars represent 1 SE. N = 6. Although error magnitudes are different from those seen in Experiment 1, the same pattern is observed as was seen in Experiment 1, namely a gradual increase in spread of errors with the number of objects tracked.

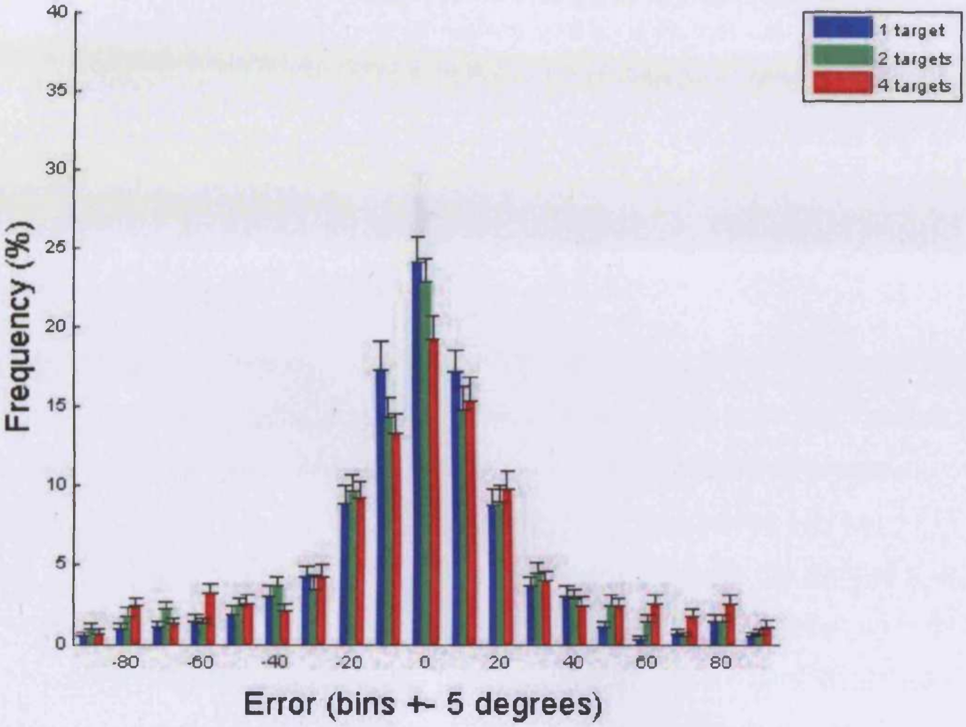

Figure 19: Errors in reported orientation for Experiment 8. This experiment used slow rates of orientation change. Error bars represent 1 SE. N = 6. Although error magnitudes are different from those seen in Experiment 1, the same pattern is observed as was seen in Experiment 1, namely a gradual increase in spread of errors with the number of objects tracked.

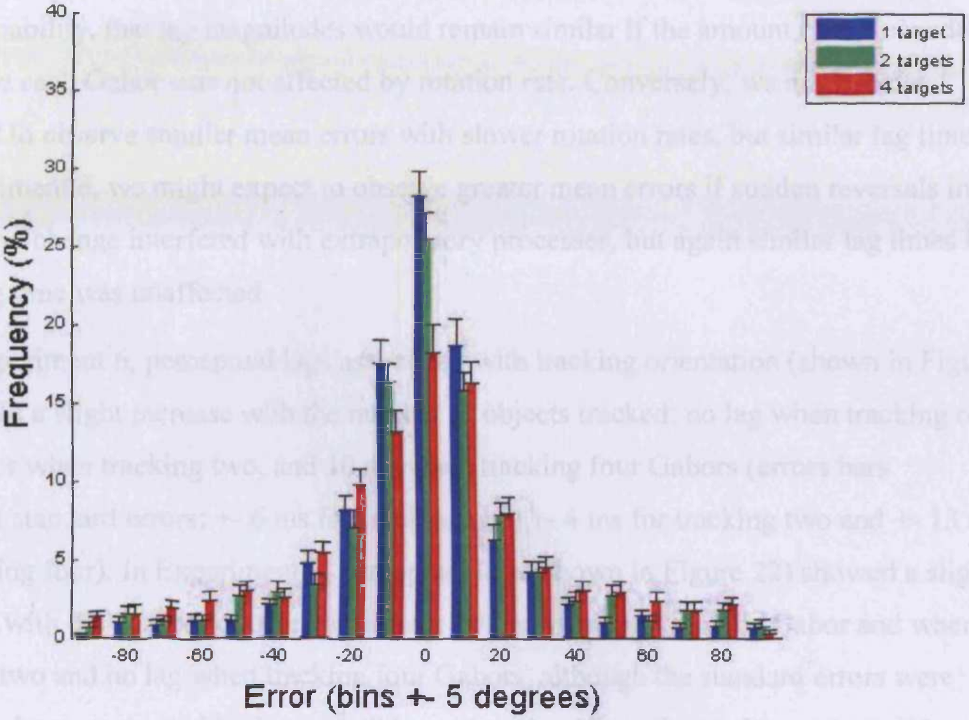

Figure 20: Errors in reported orientation for Experiment 9. This experiment used fast rates of orientation change. Error bars represent 1 SE. N = 6. Although error magnitudes are different from those seen in Experiment 1, the same pattern is observed as was seen in Experiment 1, namely a gradual increase in spread of errors with the number of objects tracked.

4.3.2 Lag analyses

The second type of analysis, and the analysis of real interest in these four experiments, relates to the perceptual lag patterns. Lags remained small in all experiments with little effect of set size and of similar magnitude to those observed in Experiment 1. Lag patterns are shown in Figures 21 – 24 below. The lags displayed by individual observers are shown in Appendix 5 for Experiments 6 - 9. Individuals display variable mean error magnitudes, set size effects and lag patterns, but none suggest the large lags and set size effect on lags that were observed in Experiment 2. It might have been expected that although mean errors might increase with faster rotation rate or with lower discriminability, that lag magnitudes would remain similar if the amount of time needed to encode each Gabor was not affected by rotation rate. Conversely, we might have expected to observe smaller mean errors with slower rotation rates, but similar lag times. In Experiment 6, we might expect to observe greater mean errors if sudden reversals in orientation change interfered with extrapolatory processes, but again similar lag times if encoding time was unaffected.

In Experiment 6, perceptual lags associated with tracking orientation (shown in Figure 21 showed a slight increase with the number of objects tracked: no lag when tracking one Gabor nor when tracking two, and 10 ms when tracking four Gabors (errors bars represent standard errors: +- 6 ms for tracking one, +- 4 ms for tracking two and +- 13 ms for tracking four). In Experiment 7, perceptual lags (shown in Figure 22) showed a slight decrease with the number of objects tracked: 10 ms when tracking one Gabor and when tracking two and no lag when tracking four Gabors, although the standard errors were relatively large compared to these small lag values (+- 19 ms for tracking one, +- 33 ms for tracking two and +- 17 ms for tracking four). Experiments 8 and 9 also showed a very slight increase in perceptual lags with the number of objects tracked (Experiment 8: 0±2 ms for tracking one Gabor, 0±20 ms for tracking two and 10±33 ms for tracking four, Experiment 9: 10±3 ms for tracking one Gabor, 10±23 ms for tracking two and 20±6 ms for tracking four).

In Experiment 6 (sharp reversals in orientation change), the correlation between the lag time and the number of objects tracked was significant (correlation coefficent = 0.446, $N = 18$, p (1 tailed) < 0.05). In Experiment 8 (slow orientation change) it was similarly significant (correlation coefficent = 0.441, $N = 18$, p (1 tailed) < 0.05). In Experiment 7 (restricted orientation range) the correlation was not significant (correlation coefficent = -0.088, $N = 18$, p (1 tailed) $= 0.364$) nor was it in Experiment 9 with fast orientation changes (correlation coefficent = 0.120, $N = 18$, p (1 tailed) $= 0.317$).

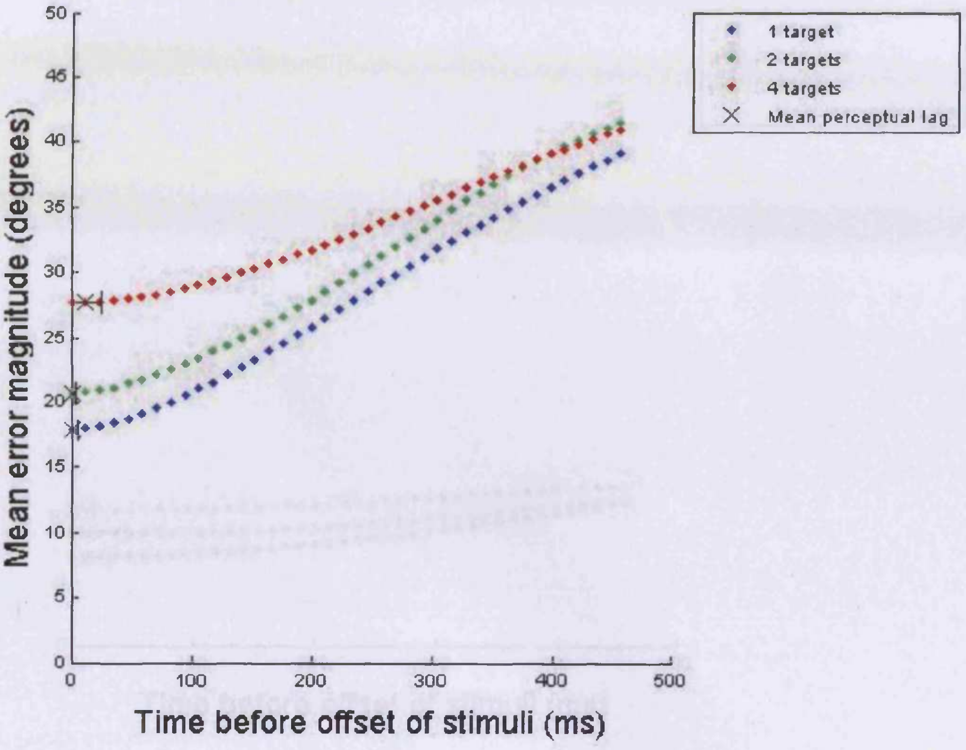

Figure 21: Lag analysis for Experiment 6. Perceptual lags are of similar magnitude to those reported in Experiment 1, despite the addition of sharp changes to orientation values. The results suggest that the large lags and set size effects on lags in Experiment 2 were not due to sharp reversals in feature change. Error bars represent 1 SE.

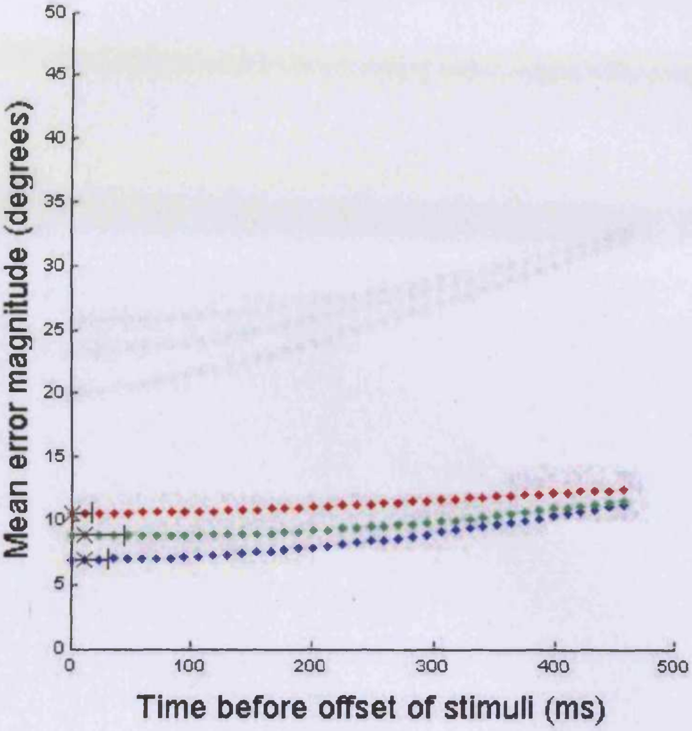

Figure 22: Lag analysis for Experiment 7. Minimum points in mean error magnitude are less pronounced, but nonetheless perceptual lags are of similar magnitude to those reported in Experiment 1, despite the restriction to the range of possible orientation values. The results suggest that the large lags and set size effects on lags in Experiment 2 were not due to discriminability of feature changes. Error bars represent 1 SE.

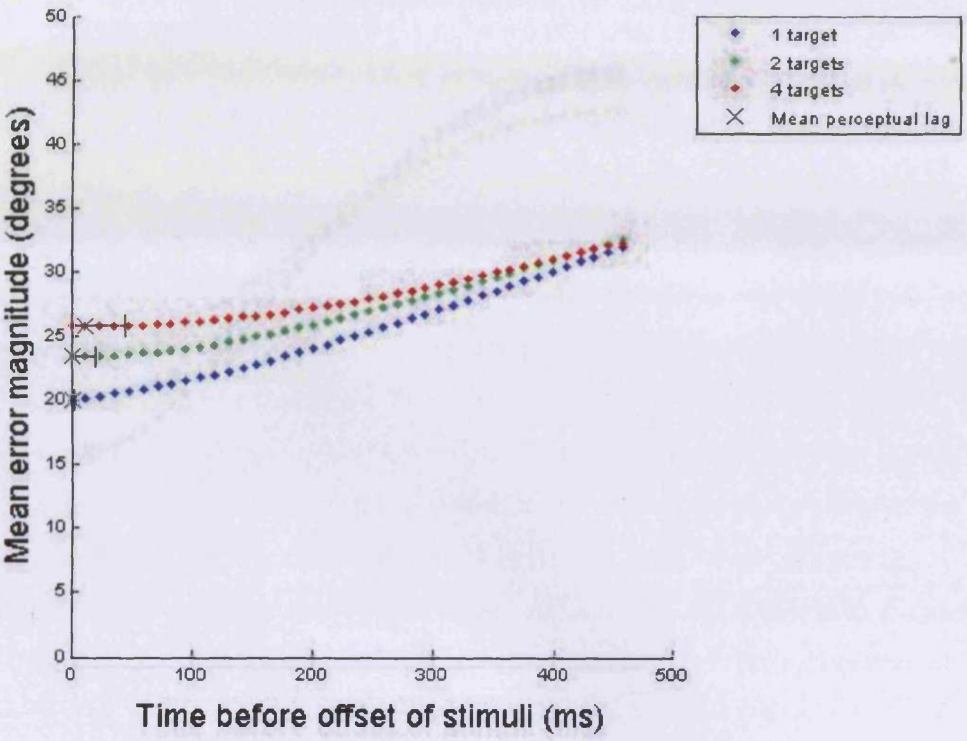

Figure 23: Lag analysis for Experiment 8. Perceptual lags are of similar magnitude to those reported in Experiment 1, despite the slow rate of orientation change. The results suggest that the large lags and set size effects on lags in Experiment 2 were not due to differences in perceived rates of feature change. Error bars represent 1 SE.

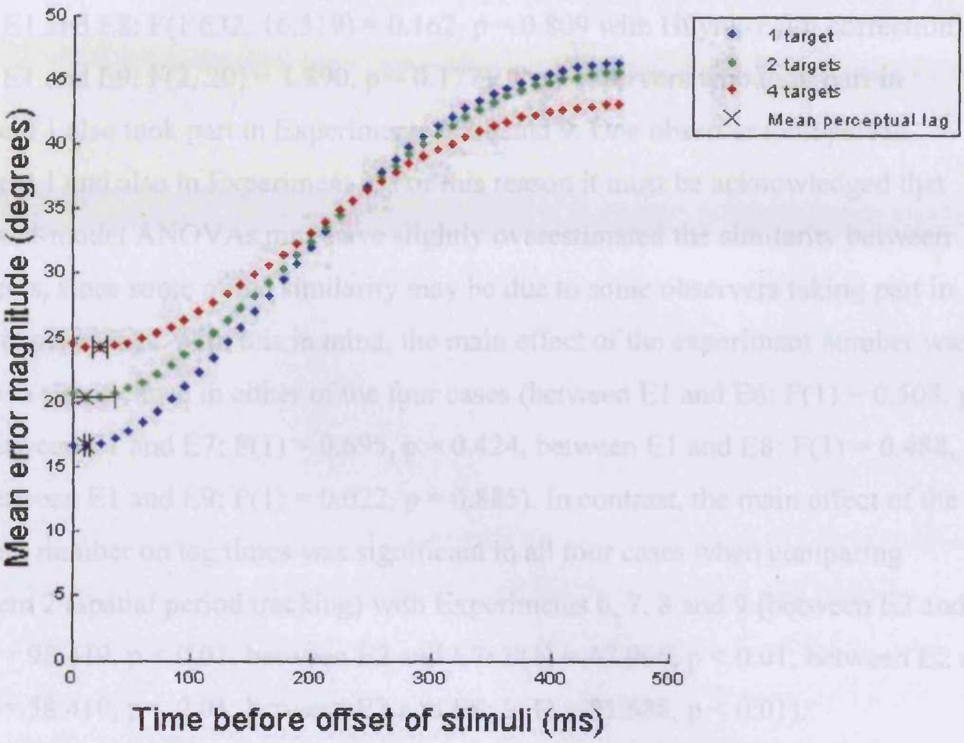

Figure 24: Lag analysis for Experiment 9. The rapid change of orientation in the stimulus yielded steeper slopes in mean error magnitudes. Perceptual lags, however, are of similar magnitude to those reported in Experiment 1 despite the rapid rate of orientation change. The results suggest that the large lags and set size effects on lags in Experiment 2 were not due to differences in perceived rates of feature change. Error bars represent 1 SE.

4.3.3 Differences in lag patterns between Experiment 1 and Experiments 6 – 9

There was no evidence for a difference between any of these four experiments and Experiment 1 in terms of the effect of set size on perceptual lag times; interactions between the effects of set size and experiment number on lag times were non-significant in 2 (Experiment number) x 3 (set size) mixed-model ANOVAs for all four cases (between E1 and E6: F(1.544, 15.438) = 0.502, p = 0.568 with Huynh-Feldt correction, between E1 and E7: F(1.337, 13.372) = 1.205, p = 0.310 with Huynh-Feldt correction, between E1 and E8: F(1.632, 16.319) = 0.162, p = 0.809 with Huynh-Feldt correction, between E1 and E9: F(2, 20) = 1.890, p = 0.177). Two observers who took part in Experiment 1 also took part in Experiments 6,7,8 and 9. One observer took part in Experiment 1 and also in Experiment 9. For this reason it must be acknowledged that these mixed-model ANOVAs may have slightly overestimated the similarity between experiments, since some of the similarity may be due to some observers taking part in multiple experiments. With this in mind, the main effect of the experiment number was not close to significance in either of the four cases (between E1 and E6: F(1) = 0.508, p = 0.492, between E1 and E7: F(1) = 0.695, p = 0.424, between E1 and E8: F(1) = 0.488, p = 0.501, between E1 and E9: F(1) = 0.022, p = 0.885). In contrast, the main effect of the experiment number on lag times was significant in all four cases when comparing Experiment 2 (spatial period tracking) with Experiments 6, 7, 8 and 9 (between E2 and E6: F(1) = 98.419, $p < 0.01$, between E2 and E7: F(1) = 47.966, $p < 0.01$, between E2 and E8: F(1) = 58.410, $p < 0.01$, between E2 and E9: F(1) = 81.688, $p < 0.01$).

4.3.4 Summary of findings in Experiments 6 – 9

None of the observed lag patterns in Experiments 6 - 9 were similar to those observed for spatial period monitoring in Experiment 2. It seems then that the difference in lag patterns between Experiments 1 and 2 was not due to the specific characteristics of the feature change, since neither sharp feature change reversals, discriminability nor rates of feature change affected the lag patterns displayed in Experiment 1. Instead, it appears that the difference in lag patterns observed between Experiments 1 and 2 is due to the different features themselves. Possible reasons why spatial period tracking exhibits different patterns from orientation and position tracking are discussed in Chapter 5.

CHAPTER 5. DISCUSSION OF CONTINUOUS MONITORING EXPERIMENTS IN CHAPTERS 2 – 4

5.1 SUMMARY OF RESULTS OF CHAPTERS 2 – 4

5.1.1 The nature of the attentional limit on continuously monitoring the features of multiple objects

The results described in Chapter 2 demonstrate a progressive decline in the precision of the representation of tracked objects, for orientation, spatial period, and location tracking. This decline occurs before any four object limit, and even when going from monitoring one object to monitoring two objects, for position and orientation monitoring. This progressive decline with the number of objects monitored is incompatible with a fixed-number-of-objects limit of around four objects, such as the FINST model (Pylyshyn, 1989) for location tracking and is also incompatible with a fixed-number-of-objects limit of around four objects for orientation or spatial period tracking. There is no way in which a fixed-number-of-objects limit of around four objects could predict the decline in performance observed in Experiments 1 – 3, instead it would predict a decline only as the number of objects monitored is increased from four objects upwards.

In all three experiments, the data are also consistent with the existence of a flexible resource of the kind discussed by Alvarez and Franconeri (2007) whereby attentional resources can be spread between attended objects according to the number attended. The tracking resource could not be entirely flexible, however, or it would be allocable in parallel to all targets, and our evidence for a serial component is not consistent with this type of complete flexibility.

5.1.2 Contribution from serial processing

For orientation, and particularly for spatial period, we find evidence for increasing perceptual lags as more objects are added to the attentional load. Analyses of error patterns broken down by guessing likelihoods suggest a contribution from serial processing, particularly for spatial period monitoring. For spatial period monitoring in particular, there are clear differences in the set size effect on lag times when responses are split into those most likely to contain a high proportion of guessing responses, and those less likely to contain a high proportion of guesses. Of course, pure guesses will exhibit no particular lag at all, but responses that draw in part from a guessing strategy can reveal lags corresponding to the time when the now relatively poor representation was last updated. That pure guesses produce no lag patterns was confirmed in the simulations described in section 2.4.7. The outcomes of these analyses are consistent with observers switching their processing resources between targets for tracking, such that lag times are at a constant and small value when the queried object is the last-processed object or objects. When the queried object is not the last-processed object, lags are increasingly longer the more objects must be attended, implicating a serial process where the lag reflects the last time on average that the object was updated.

Further support for the notion that the lags are caused by serial switching among stimuli was suggested by the results of Experiments 4 and 5. In these experiments, observers first reported one target in exactly the same way as in Experiments 1 and 2 but then were immediately cued again to report a second target. We inferred that serial and parallel processing should cause different patterns in the correlations between accuracy in reporting a first-queried and a second-queried target. Serial processing should show up as a negative correlation, because the two reports would compete for the same attentional resource. Parallel processing, however, should show up as a positive correlation, because both reports would be affected by the same level of availability of the attentional resource, for instance as affected by arousal on that trial. Accuracies in reporting both of two monitored objects were correlated for orientation but not spatial period, suggesting of a greater contribution from serial processing in the case of spatial period than for orientation.

More serial processing can then explain the greater effect of set size on perceptual lag for spatial period than for orientation. More serial processing would cause an increase in the effect of set size on lags, since if resources are being switched between objects, the more objects are attended, the longer in the past a queried object would have last been processed, on average. Of course, if an observer is serially processing four objects, say, then on one quarter of trials, one would expect the queried object to have been the last attended object and thus to exhibit a relatively minimal lag. However, the other 75 per cent of trials will exhibit a range of greater lags, since the queried object will sometimes have been potentially attended as far in the past as it takes for attention to have switched through all of the objects.

Experiments 6 - 9 rule out the possibilities that the observed differences in perceptual lags between orientation and spatial period were due to differences in discriminability, rates of feature change or sudden reversals in feature change. Experiments 6 – 9 involved modified versions of the range of orientations used in Experiment 1 to simulate the feature properties of the spatial period range used in Experiment 2. None of the tested modifications of the feature range induced the set size effects in lags that were observed for spatial period monitoring. This it seems unlikely that the lag patterns observed in Experiment 2 were due simply to the range of possible spatial periods selected for the stimuli.

5.2 ADVANTAGES OF THE CONTINUOUS MONITORING PARADIGM

The data here provide a measure of the internal representations of attended objects through the use of adjustable sample stimuli and are free of factors that may have inflated the set size effect observed in previous studies. Such additional factors include the comparison processes required for change detection tasks and memory requirements of VSTM tasks. The data are also free of contamination by the type of decision noise discussed by Palmer (1995) whereby performance will decline with increases in set size even in a purely parallel process with infinite capacity, as the chance of errors associated with the representation of one item will be replicated with each addition to the set size. Here there is no role for this type of decision noise since responses here were based on the representation of one item only.

5.3 ICONIC MEMORY AND VISUAL PERSISTENCE

To perform the continuous monitoring task, one might have expected that observers would tap into a classic iconic memory store or visual persistence that decays rapidly after stimulus offset (Averbach & Coriell, 1961; Coltheart, Lea & Thompson 1974; Sperling, 1960). As mentioned in the introduction, the classic conception of a high-fidelity iconic memory might suggest no decrement in precision with the number of items monitored, as long as items are immediately queried after disappearance and as long as only one of the monitored items is ever queried. These are exactly the conditions of the studies reported in Experiments 1 – 3 and hence we might expect observers to simply access the last feature value of the queried object directly from high-capacity iconic memory. If iconic memory were determining performance, the addition of a short delay (270 ms) between display offset and post-cue should substantially reduce performance (e.g. Lu, Neuse, Madigan, & Dosher, 2005) but in pilot tests it did not. Thus observers apparently are not able to use iconic memory after display offset, and the reason for this is not clear. It could conceivably be due to the continuously changing nature of the stimuli, or the reporting of values along a feature dimension, which differ from classic iconic memory studies.

Furthermore, if after stimulus offset, observers were simply able to use high-capacity iconic memory to report the last feature value as described above, based on the classic conception of iconic memory, we would expect no change in lag with set size. With negligible retention interval and only one object to report, regardless of set size, observers might immediately direct their attention to the appropriate object and retrieve the representation stored in iconic memory before it decays. Of course, there may actually be lags in the iconic memory representation, it is just that there was no way in which a perceptual lag could have revealed itself in classic tasks such as letter reporting. Although perceptual lags have not been explored explicitly in an iconic memory paradigm, Smithson and Mollon (2006) have recently suggested that iconic memory might include a temporal dimension, such that masks need not necessarily overwrite previously presented target stimuli, for instance. Rather, the target and the mask could both reside in the iconic store but at different points along the temporal dimension. In a task where overwriting previously presented images is advantageous, such as in Chapters 2 – 4, it seems that previous states of changing objects might persist in the iconic representation alongside most recent states, and potentially could mean that lags might at least partially reside in the iconic store.

Whilst observers do not appear to be accessing high-capacity iconic memory after offset of the stimuli, it is possible that mnemonic processes are involved during the monitoring task while the stimuli are still being displayed. The capacity limit evidenced by the increase in lags with the number of objects monitored could be incurred at a perceptual stage or a mnemonic stage. Either perception or loading into memory, or both, could involve serial or slowed parallel processing with increases in set size.

5.4 RELATION TO PREVIOUS POSITIONAL FEATURE TRACKING STUDIES

Like the results of Experiment 3, Tripathy and Barrett (Tripathy & Barrett, 2003a; Tripathy & Barrett, 2003b; Tripathy & Barrett, 2004) also found evidence for a decline in spatial tracking performance with set size. Their task required observers to track moving objects, one of which would undergo a deviation in its trajectory and observers then reported the direction of this deviation. Theirs was a different type of task, however, as observers had to monitor for a change in direction of motion which is likely to have involved processes other than monitoring the positions of objects. Observers may have been monitoring the directions of motion, or they may have used two or more points along a trajectory to work out the likelihood of a trajectory having changed, for instance.

When the task requirement was specifically to monitor positions, we show a progressive decline in the precision of the location representation of tracked objects. However this decline may not be as large as that found by Tripathy and colleagues. They found that observers could reliably discriminate trajectory deviations of 19 deg or greater when tracking two objects, but for four or five objects a deviation of 76 was needed (Tripathy, Narasimhan, & Barrett, 2007). This nearly four-fold increase in imprecision is much greater than the approximate doubling in mean position error that we found. However, as we have mentioned, detecting a trajectory deviation may have different demands and further work bridging the two tasks is needed.

The positional imprecision reported here is likely to contribute to the observed set size effects measured by others in MOT tasks (e.g. Alvarez & Cavanagh, 2005; Alvarez & Franconeri, 2007; Alvarez et al., 2005; Keane & Pylyshyn, 2006; Pylyshyn, 1989; Pylyshyn & Storm, 1988; Scholl & Pylyshyn, 1999; Scholl et al., 2001; vanMarle & Scholl, 2003; Yantis, 1992). We ran preliminary simulations of predicted performance on a standard MOT task using the imprecision errors recorded here to predict how often observers would 'lose track' of targets by confusion with non-targets.[3] Positional imprecision could cause observers to erroneously switch from tracking a target to a non-target when objects get close to one another. From the positional errors recorded in Experiment 3, we were able to simulate how often positional uncertainty could cause a target and a non-target to be confused with one another, and how this would become increasingly probable the closer the two objects become. We found that simulated performance was much worse than is typically observed in standard MOT tasks. This points towards the use in MOT tasks of strategies such as speed or velocity (speed and direction) monitoring to recover lost targets for tracking and to disambiguate targets from non-targets. Indeed, a similar pattern of a gradual decrease in sensitivity for the directions of motion of tracked objects has recently been reported by Horowitz, Place, Van Wert and Fencsik (2007) as the number of MOT targets is increased. Of course, it may also be the case that positional imprecision errors are even greater when observers are also

---

[3] Details of these simulations can be found in Appendix 6.

deliberately attempting to track trajectory characteristics. Equally, it could also be the case that precision is better in MOT tasks than is tapped by explicit report.

In either case, regardless of whether observers undertake trajectory monitoring as well as position monitoring in MOT tasks, and whether positional precision is either better or worse than our estimates, our spatial imprecision result in Experiment 3 supports the conclusions recently reached by Alvarez and Franconeri (2007). They found that performance in MOT tasks is limited by the spatial resolution of attention, which is in turn dependent on attentional load. Although other processes, such as direction tracking, for instance, may be at work in MOT tasks, the spatial imprecision we observe in Experiment 3 appears to be a likely candidate for at least some of the observed set size effects.

Our spatial precision capacity result can throw some light on the recent findings of Franconeri, Alvarez and Enns (2007) who found that the estimated capacity limits for spatial selection depended on the level of precision required. After precues to a variable number of potential target locations, observers performed a visual search on the basis of objects' orientations. Capacity estimates were dependent on the level of precision necessitated by the density of objects in the display. As Franconeri et al. discussed, their observed capacity estimates may have indicated either a limit on spatial selection or on subsequent encoding of orientations. Decreasing precision with attention to increasing numbers of objects explains their results even without considerations of the capacity limit for attending to multiple orientations. This interpretation is also consistent with evidence that attention acts to increase spatial resolution at an attended location for the purpose of texture segregation (Yeshurun & Carrasco, 1998) and localising a briefly presented object (Tsal & Bareket, 1999).

5.5 RELATION TO PREVIOUS RESULTS RELATING ATTENTION TO PRECISION

Our observation of a progressive decline in the precision of representations of tracked objects is consistent with the findings of Wilken and Ma (2004) who found a gradual decline in precision in visual short-term memory tasks, rather than continuous attentional tracking. Their results were similar to those in Experiments 1 -3 in that they showed an increase in the spread of errors made in reporting the features (either the colour, orientation, or spatial frequency) of multiple objects, with increases in set size. This is not surprising if attentional capacity is a processing bottleneck before information can reach VSTM (e.g. Cowan, 2000; Duncan & Humphreys, 1989; Duncan & Humphreys, 1992), nor indeed if attention is gating access to encoding into VSTM where this measured capacity limit could reside.

Prinzmetal, Amiri, Allen and Edwards (1998) carried out studies similar to Experiments 1 – 3, without an explicit memory component to the task, but using a dual-task instead of cueing to manipulate attention. Prinzmetal et al. showed that attention reduced the variability of responses when observers reported the features of briefly presented objects either in the presence or absence of a secondary task. We are to demonstrate this finding here in terms of the gradual decline in precision with the number of objects attended, as in Prinzmetal et al.'s task there was no systematic exploration of the effect of attending to multiple objects. In addition our method presented here is free of any possible confounding effects brought about through the use of a dual task. Their secondary task was a visual search amongst an array of letters, which presumably involves a number of processes such as letter recognition, holding the target in memory, moving attention between letter-locations and marking previously searched locations. Hence a secondary task of this type may involve a number of processes, the effects of which may be difficult to predict on performance on the primary task.

The effects on the precision of representation we observe in Experiments 1-3 reflect the signal-to-noise ratio, or in the terms of signal detection theory, the parameter d'. In a dichotic listening task, Broadbent and Gregory (1963) discuss the independence of this measure from decision criteria, or the likelihood of making a 'target present' response at a given level of signal-to-noise. They assessed the effects of attention on d' by presenting listeners with burst of noise to one ear, which may or may not contain a target tone. The other ear was presented with six digits. Listeners had to report the degree of confidence that the tone was present, either when they were required to report or to ignore the six digits. Similarly to our conclusions in Chapter 2, they report that the effect of attention, as manipulated through the secondary digit reporting task, was to increase or decrease d', the signal-to-noise ratio.

5.6 PERCEPTUAL LAGS

As well as a progressive decline in the precision of the representation of tracked objects with increasing set size, we here show an increase in perceptual lag, particularly in the case of spatial period (orientation, E1: tracking one target – 0 ms, tracking two – 10 ms, tracking four – 40 ms; spatial period, E2: tracking one target – 140 ms, tracking two – 210 ms, tracking four – 250 ms). There is also a non-significant trend for an increase in lags with the number of objects tracked for position. In other words, there was a tendency for observers to report features from the past, and the more objects were attended, the greater this lag became.

Carlson, VanRullen, Hogendoorn, Verstraten and Cavanagh (2007) have observed lags of similar magnitude for a task similar to the orientation monitoring presented here. They asked observers to monitor between one and five 'clock-like' stimuli simultaneously. They observed a small lag associated with each addition to the number of objects monitored of about 12 ms, up to an asymptote of around a total lag of 100 ms for four clocks. There are two reasons, however, why their observations may not represent lags fairly. Firstly, their stimuli were similar to analogue clock with a single 'hand'. Different 'hand' readings were represented not only by an orientation but by a different spatial location on the 'clockface'. Our orientation monitoring task here is much more difficult to perform on the basis of location alone since the Gabors occupied constant positions. Secondly, the motion of the 'hands' was entirely predictable as each had a constant angular velocity. This leaves open the possibility that observers were able to extrapolate readings based on the reading when that particular clock was last attended. For these reasons, the lags observed for orientation monitoring here may reflect a 'purer' measure of people's ability to continuously monitor the orientations of multiple objects.

It is possible that the perceptual lags observed in Experiments 1 – 3 could result to some extent from the sensory integration process. If it takes time for information to accumulate in the visual system in order to produce a reportable percept, then we would of course expect to find non-zero lags for the case of monitoring one object. This seems reasonable since, for instance, cortical cells sensitive to spatial frequency exhibit dependence on the temporal duration of stimulus presentation (e.g. Movshon, Thompson & Tolhurst, 1978). If, for instance, a task involves a sensory integration period of x ms, then we would expect observers to report a value close to the mean feature that the object possessed during this period, corresponding to the feature value of the object x/2 ms in the past. Sensory integration cannot, in itself however, explain the increase in lags with the number of objects monitored, without assumptions about the dependence of this integration process on the number of objects being processed.

The reasons for finding such large perceptual lags in the case of spatial frequency monitoring for one object are not clear. This could potentially be due to the sustained demands of the continuous monitoring task. Alternatively this lag could reside in the process of continuously updating one's representation in memory of the monitored object.

One might also expect that the continuously changing nature of the stimuli could contribute to the observed lags reported in the case of monitoring one object, if Gabors presented on subsequent frames produce either forward (e.g. Schiller, 1966) or backward (e.g. Breitmeyer & Ogmen, 2000; Ogmen Breitmeyer & Melvin, 2003) masking, or both. As far as we are aware, there are no reports of the effects of such masks in the literature i.e. when the mask and target are effectively the same continuously changing stimulus. However, as we previously mentioned, Smithson and Mollon (2006) suggested that masks need not necessarily overwrite the representations in iconic memory of previously presented stimuli. They suggest that iconic memory might include a temporal dimension, and it appears possible that in a task like continuous monitoring, where many successive stimuli are presented from frame to frame, this temporal dimension acts to store both present and past states of the continuously changing object.

5.6.1 Relation of perceptual lags to other phenomena in motion perception

The flash-lag effect (e.g. Nijhawan, 1994) is a well-documented and tangentially related temporal phenomenon, whereby a briefly presented ('flashed') object is displayed at the moment when a continuously moving object passes close by, and the flashed object appears to spatially lag behind the moving object. It has been additionally demonstrated to apply to the continuously changing features of colour, luminance and spatial frequency (Sheth, Nijhawan & Shimojo, 2000) as well as to the perceived positions of objects, However, this literature is only tangentially related to the studies presented here in Chapters 2 - 4, since it applies to the relative spatial and temporal perceptions of two disparate (moving and stationary) objects.

More closely related is the phenomenon of representational momentum (e.g. Freyd & Finke, 1984; Thornton & Hubbard, 2002), or the tendency to remember the last position of moving objects at a position displaced in the direction of its motion. That we do not observe this tendency in Experiment 3 may be explained by the role of eye movements in representational momentum studies. Tracking eye movements, when combined with visible persistence, have been suggested to contribute to forwards-mislocation of moving objects (Kerzel, 2000; Kerzel, 2006).

Our findings with respect to a lack of representational momentum in MOT are consistent with those of Keane and Pylyshyn (2006). In their task, there were multiple moving targets for tracking that would be difficult to track simultaneously with eye movements. Hence under Kerzel's account, we should not expect representational momentum. They found that if objects disappeared briefly during an MOT task, observers were better when objects reappeared at their last positions or at recent previous positions, than when they reappeared at extrapolated locations. The results of our Experiment 3 suggest that perceptual lags may have contributed to the Keane and Pylyshyn result.

Indeed our lag result in Experiment 3 is also consistent with the findings of Roulston, Self and Zeki (2006) who observed a small "reverse representational momentum" effect. They showed that two bars moving towards one another had to move slightly beyond one another in order to be reported as being spatially aligned at the moment of stimulus offset. This absence of representational momentum can be explained since Roulston et al. presented their stimuli foveally, and since Kerzel, (2000, 2006) suggested that representational momentum requires tracking eye movements combined with the effects of visible persistence. The contribution in Roulston et al.'s study made by the relative judgement between two moving objects is not clear. For instance, comparison processes between stimuli may have affected their results, and our result in Experiment 3 is free of such concerns.

5.6.2 Implications of perceptual lags

This finding of greater perceptual lags when more objects are monitored may have significant repercussions for everyday tasks such as navigating opponents on a football field or simply avoiding vehicles when crossing the street. In the psychophysical literature, perceptual lag is commonly measured in the flash-lag effect but investigation of the number of objects monitored seems to be limited to objects on a common path which would perhaps be grouped together (Khurana, Watanabe, & Nijhawan, 2000). Our methodology might be adapted to test for a serial or limited-capacity component to the flash-lag phenomenon. Future work should also test for the possible consequences in everyday actions, such as making eye movements towards or reaching to intercept one of a number of moving objects.

5.7 PARALLEL AND SERIAL PROCESSING

In the tasks presented here, observers attempted to attend to all the targets at once. Despite this, we still find evidence for serial processing of to-be-attended objects in the tasks. The pattern of responses observed, particularly for spatial period tracking in Experiments 2 and 5 suggests a contribution of serial processing between attended objects when tracking spatial period values. Experiments 2 and 5 showed large lags for spatial period tracking and a large effect of the number of objects tracked on lags. In addition, the pattern of lags broken down by guessing likelihood was consistent with a serial element to processing. Responses less likely to have drawn from a guessing strategy were associated with no set size effect, and responses more likely to have drawn, at least in part, from a guessing strategy exhibited an exaggerated set size effect. This pattern suggests that responses were more accurate precisely on those occasions when the queried objects was the last object attended, pointing towards switching of processing between objects. Furthermore, the lack of correlation between first and second reports in the double report experiment (E5) is consistent with contributions from parallel and serial processing. Although this was a null result, and therefore only weakly informative of the underlying relationship, it was logically consistent with the presence of a serial element to processing of spatial periods. Of course, switching could reflect either serial attentional or mnemonic processing. In either case, however, it would reflect a serial component to the processing of objects.

Recently, others have found evidence for serial processing of multiple stimuli when stimuli were relatively difficult to discriminate but not when they were relatively easy to discriminate. VanRullen, Carlson and Cavanagh (2007) attempted to resolve the issue of whether multiple objects are attended through serial or parallel processing. They measured the psychometric functions for detection of a contrast decrement as a function of presentation time, and used this to model predicted functions from three possible attentional mechanisms. The three models were a ‘parallel’, a ‘sample always’ and a ‘sample when divided’ model: in the first, attention is a limited-capacity resource but is spread simultaneously between targets, in the second, attention is shared between targets by successive sampling, but is always a periodic process even when attending to one object, and in the last, attention samples periodically, but only when attending to more than one object. They found that the best fitting model depended on the difficulty of the detection task: for the more detectable stimuli, observers appeared to be using parallel processing, but for the more difficult version they appeared to be sampling at a rate of around seven samples per second, even when attending to only one object. VanRullen et al. likened this to a ‘blinking attentional spotlight. However, that they did not find evidence for either of the serial models for the more detectable stimuli may have been due to the type of detection task used. Observers knew that there would be a maximum of one contrast decrement event associated with one of the four objects, on any one trial. It seems possible that the onset of the contrast decrement on one of the objects was sufficient to prevent switching of attention away from this object, thus reducing any serial processing observed. Our paradigm is free of this type of concern, since features on all the target objects in Chapters 2 – 4 were entirely independent at all times.

Whilst VanRullen, Carlson and Cavanagh (2007) find limited evidence for a serial component to processing of multiple stimuli, Carlson, VanRullen, Hogendoorn, Verstraten and Cavanagh (2007) claim to find evidence for purely parallel processing. However, their study confounded orientation and position tracking, and also allowed the possibility of observers extrapolating future states of the stimuli. Carlson et al. (2007, Experiment 2) report a similar double-report method and inter-error correlations to those we present in Experiments 4 and 5 for monitoring two of their 'clock-like' stimuli. They asked observers to monitor two clocks, and to report both readings. A purely serial model of processing would predict a negative correlation between accuracies of the two reports, for the reasons explained in Experiments 4 and 5 here, namely that under serial processing, full attention to one object necessarily entails that less attention will be directed to the other. Carlson et al. observed no correlation between the two reports, and interpreted this as evidence for a limited capacity parallel model. It is, however, also consistent with a combination of both serial and parallel processing: purely parallel processing would result in a positive correlation since small fluctuations in processing resources would affect both reports in the same fashion. Again, however, because of the changing orientations and positions of the hands on their 'clock-like' stimuli, their results most probably reflect a combination of position and orientation monitoring. Their data are vulnerable to inflation of measured capacity because of the possibility of extrapolating clock readings based on their constant angular velocities. In the studies presented in Chapters 2 -4, extrapolation would be difficult because the accelerations and velocities of the stimuli through features were both constantly variable.

Bichot, Cave and Pashler (1999) attempted to demonstrate that attention can be deployed simultaneously to the colours or orientations of two objects at two different locations. They based this conclusion on the fact that observers were no better at responding to two colours (or two orientations) when they were presented successively than when they were presented simultaneously. However, their simultaneous presentation condition allowed observers to view the stimuli for around 150 ms, which according to our lag results reported here, is almost certainly long enough for some serial processing of orientation, at least.

Recently, Huang and Pashler have referred to the results presented here in Chapter 2 as being consistent with their Boolean Map Theory (Huang & Pashler, 2007). Boolean Map Theory is used to describe the mechanism of attention in selecting various objects on the basis of their features and through top-down attentional control. It claims that we can only have perceptual access to one value of each feature at any instant in time, for instance, to select one particular colour or one orientation out of many possible orientations. In their theory, there is no cost associated with selecting multiple objects with this feature value, and each object with that value is linked to its spatial location. This means that in selecting a feature value, we necessarily have simultaneous access to a 'map' of all locations at which objects possess this feature value. This causes an asymmetry between location-based and featural information, as only one feature value can be accessed at any one time, but multiple spatial locations can be accessed. It also includes the prediction that performance costs will be associated with selecting multiple feature values, for instance attending to both blue and red objects in a scene. Evidence for the theory was provided by Huang, Treisman and Pashler (2007). They presented two stimuli either simultaneously or in successive presentations. They then queried observers' perceptions of the objects' location or colour. They found an interaction between the effects of the timing of presentations and whether colour or location was queried. For reporting locations, observers performed better with simultaneous presentation, but for reporting colours, they performed better with the successive presentations. In a second experiment, the two objects were always presented simultaneously. In one condition, observers did not know which objects would later be queried, and in the other they were told this in advance. Prior knowledge had a greater effect on reporting colour than location, suggesting that directed attention was required to access the colour of each differently coloured object, in contrast to the immediately accessible locations of both objects. The results presented in Chapters 2 - 4 are consistent with their theory in that we find evidence for a likely serial component when monitoring the spatial periods of multiple objects. This might be predicted from Huang and Pashler's theory since they explicitly predict that multiple values of the same feature cannot be easily selected simultaneously. In contrast they predict no cost associated with selecting multiple positions simultaneously, in accordance with the much smaller lags we observe for monitoring positions. That orientation tracking appears to be more similar to position tracking than to spatial period tracking might be explained by possible encoding of the orientation stimuli in terms of positions, as discussed below in section 5.8.

It seems possible that parallel and serial processing might produce different patterns of performance due to different attentional processes: for instance, Posner and Petersen (1990) have proposed several processes in the spatial allocation of attention that are distinct from processing of visual information once attention has reached a target's location. In Posner's classic task (1980), observers fixate a central point and a cue indicates that a target stimulus will appear either to the left or the right of fixation. The cue can either be valid (in that it correctly indicated where the target would appear), neutral or invalid (indicating the opposite direction to where the target would subsequently appear. Observers are faster to react to validly cued targets than trials containing a neutral or invalid cue. Additionally, reaction times are longer after an invalid cue than a neutral or valid cue. Posner and Petersen interpreted this as evidence for the direction of attention to the location of the target on valid trials, and to the invalid location on invalid trials, thus causing a difference in reaction times to the target. Furthermore, they propose three processes involved in the spatial allocation of attention: a process of disengagement from the currently attended location, a spatial shift of attention and subsequent re-engagement of attention at the new location. Only once these processes have taken place can processing of a target at the new location begin.

The validity of these distinctions has since been supported by many studies. For instance, Snowden, Wiley and Muir (2001) found a manipulation of target contrast to primarily affect the movement of attention to targets, rather than subsequent processing of the target. In special populations, specific deficits amongst these types are thought to be present, for instance, that in a classic Posner task, older observers are particularly slowed in invalidly cued trials (Nissen & Corkin, 1985) which is likely to be due to problems disengaging from currently attended locations.

If parallel and serial processing are both present to differing extents in the continuous monitoring experiments presented here in Chapter 2, one might ask at which of these attentional stages the parallel or serial elements reside? Indeed this appears to be an open question and deserving of further research. For instance, evidence for a serial element could indicate a relatively long period required to process targets once attention has moved to their location (under this view one might consider parallel processing to be a special case of serial processing with a negligible dwell time and very fast switching rate between objects), or it could indicate that attention is more frequently required to move between targets at different locations (if fewer objects can be simultaneously 'spotlighted' by attention).

5.8 DIFFERENCES BETWEEN FEATURES

We find substantially different patterns of results for orientation, spatial period and position. Three pieces of evidence all point to the serial processing component being greatest when monitoring spatial periods. The lags were greatest by far in the case of spatial period, the correlation between errors was smallest, and the analysis of responses by guessing-likelihood also yielded the largest difference.

Why is spatial period tracking so serial? Unfortunately, no direct evidence on this point appears to be present in the literature so we must enter the domain of speculation based on results from further afield. The dorsal processing stream of the brain, projecting to the posterior parietal lobe, has been implicated in the MOT task (Culham, Brandt, Cavanagh, Kanwisher, Dale & Tootell, 1998; Jovicich, Peters, Koch, Braun, Chang & Ernst, 2001). This is compatible with the dual streams of cortical processing relating to the 'what' and 'where' of stimuli in the ventral and dorsal pathways respectively (Mishkin, Ungerleider & Macko, 1983; Ungerleider & Mishkin, 1982), or more recently the distinction between dorsal 'vision for action' and ventral 'vision for perception' (Milner & Goodale, 1995).

We might expect both position and orientation tracking to be processed by the dorsal stream for two reasons. Firstly, Perenin and Vighetto (1988) showed that individuals with parietal lobe damage had difficulty in making visually guided hand movements that involved both grasping at the correct location, and orienting the hand appropriately with respect to a target object. Secondly, it is logically possible that in an orientation tracking task, orientation may itself be encoded by tracking the locations occluded or occupied by the oriented stimulus at any one time. Hence, we might expect orientation tracking to exhibit similar characteristics to location tracking if both processes are subserved by the dorsal processing stream. Conversely, we might expect spatial period to be processed by the ventral stream as it relates to pattern characteristics of objects. Additionally, in the task presented here, spatial period may not easily be encoded by tracking the location of any of the bars in the Gabor patch, for instance, because the phase of each of the five Gabors was randomised on each trial, as was that of the sample Gabor patch. Knowing the location of any of the bars in the queried Gabor would not help in adjusting the spatial period of the sample patch. For these reasons it may be reasonable to expect different patterns of results for spatial period tracking on one hand, and location tracking or orientation tracking on the other. In any case, these results demonstrate that tracking the features of multiple objects may involve different patterns of processing depending on the feature being tracked.

5.9 FUTURE DIRECTIONS

The experiments described here in Chapters 2 - 4 concern the nature of the capacity limit for simultaneous attention to the features of objects; their orientations, their spatial periods are their spatial positions. If, as others have suggested, (e.g. Magnussen & Greenlee, 1997; Huang & Pashler, 2007) it is only different instances of the *same* feature (e.g. multiple colours, multiple orientations etc.) that compete with one another, then we would expect to find no performance decrement in our continuous monitoring paradigm if observers continuously monitored both orientation *and* spatial period, for instance. If Huang and Pashler (2007) are correct in their assertion that features are encoded along with compulsory encoding of the locations of those feature values, then we would expect to observe no performance decrement for continuously monitoring non-positional features as well as their locations, simultaneously. Of course, we might expect to observe some kind of performance decrement associated with making two responses, for instance a decrement associated with holding one response in memory whilst making the first response. However, these would not be performance costs arising from any attentional limits on performance. This would certainly be a prediction worth testing using this paradigm.

It would also be interesting to compare more directly the results obtained in these experiments with those of Blaser, Pylyshyn, & Holcombe (2000). Their task was very similar to the one just described i.e. monitoring more than one feature either all from the same object, or from two spatially superposed objects. Their experiments, however, contained a slightly different task from simply reporting one of the features. Instead they had observers report the direction of a discontinuity in change along one of the feature dimensions. They observed a capacity limit of what looked like one object, although the contribution of spatial superposition of the two objects was not clear. To determine whether indeed there was some kind of competition between objects that was specific to their superposition, a similar experiment could be done with spatially separated objects.

CHAPTER 6. EXPERIMENTS 10 AND 11: COMPARISON OF ATTENTION AND CROWDING: EFFECTS ON NOISE AND SAMPLING

6.1 INTRODUCTION

In Chapter 6, we investigate the findings of Dakin, Bex, Cass, and Watt (2007) that attention does not act to decrease noise in encoding the orientations of stimuli. They instead claimed that attention acts to increase sampling of information from stimuli. We compare their methodology, a dual task attentional manipulation, with a spatial cueing methodology. We then compare the effects of attention with the effect of crowding, which was shown by Dakin et al. to increase the noisiness of encoding.

6.1.1 Feature integration and origins of attentional explanations of crowding

In 1980, Treisman and Gelade proposed a feature-integration theory of attention. The theory describes attention as the very process through which the various features of objects are bound together, resulting in coherent representations of objects. As will be explained, this view of attention has led some to investigate attentional explanations of crowding.

In visual search tasks, observers must search for a particular target among a number of distracter objects, and the target can usually be either present or absent in the display. A critical finding from such tasks is that typical search times are dependent on the way that the target is defined. In a single feature search, the target is defined by the presence of one feature, red, for instance. In a disjunction search, a target is defined by its fulfilment of either of a pair of possible features, red *or* leftward tilt, for instance. In a conjunction search, it is defined by the combination of a pair of features, red *and* leftward tilt, for instance. Search times are then plotted against the number of distracter objects to derive a search slope measure of performance.

Treisman and her colleagues (Treisman & Gelade,1980; Treisman & Gormican, 1988) observed an apparent difference between serial and parallel searches. In their studies, searches defined by a conjunction of features were associated with much steeper search slopes than either disjunction or single feature searches, to the extent that disjunction and single feature searches were relatively unaffected by the number of distracters. The target appeared to 'pop-out' from the background. Conjunction searches, in contrast, appeared to involve a slow and effortful search involving the serial redirection of attention to objects in the array before the target could be identified.

This evidence was used to derive Treisman and Gelade's (1980) feature-integration theory of attention in which this pattern of data is explained with reference to feature binding. In feature-integration theory, it is attention that allows individual features to be bound together into coherent object representations, allowing conjunctions to be identified. An important aspect of this theory is that attention is only required at the stage of whole object perception, as all individual features are available to perception pre-attentively, which is what allows 'pop-out' to occur.

Feature-integration theory cannot, in itself, explain all the data from visual search experiments as other factors have also been shown to greatly influence search slopes. For instance, Duncan and Humphreys (1989) showed that the degree of similarity between targets and distracters contributed to search times, as did the degree of dissimilarity between distracters themselves. Additionally, others have questioned whether, in fact, conjunction searches do involve a strictly serial search (McElree & Carrasco, 1999). Despite these more recent criticisms of feature integration theory, two ideas about attention persist in the literature: that not 'paying attention' to an object necessarily entails that there will be uncertainty as to which features actually belong to that object, and which are possessed by nearby objects. Similarly, that if attention cannot select an area small enough to cover one object only, then there may also arise this kind of uncertainty. For these reasons, the feature-integration view of attention has driven some to propose attentional explanations for the phenomenon of crowding.

6.1.2 Support for an attentional explanation of crowding

Crowding is the phenomenon whereby identification of the features of an object becomes difficult in the presence of proximal 'distracter' objects (Bouma, 1970; Bouma, 1973; Toet & Levi, 1992). Some have suggested that crowding occurs on those occasions when attention fails to perform its apparent feature-integration role (Treisman & Gelade, 1980) appropriately i.e. when attention binds features over an inappropriately large area. Even without endorsing feature integration theory as such, one can claim that crowding occurs when attentional resolution is not fine enough to select information from the target object only.

Attentional explanations (e.g. He, Cavanagh & Intriligator, 1996; Intriligator & Cavanagh, 2001; Tripathy & Cavanagh, 2002) place the locus of crowding at a higher level of processing than at primary visual cortex. He, Cavanagh & Intriligator (1996), for instance, observed adaptation to the orientation of stimuli that were rendered imperceptible by crowding. In addition, they observed that the crowding effect was stronger in the upper visual field, a difference not reflected in primary visual cortex. They interpreted these results as pointing towards a late locus of crowding at the stage of attentional selection. More data suggested to be consistent with an attentional account was provided by Tripathy and Cavanagh (2002) who showed that crowding was not dependent on the size of targets or the distance between the nearest edges of targets and distracters, but rather their centre-to-centre separation. This does not seem consistent with accounts based on lateral inhibition or summation at the level of simple or complex cells, for instance (e.g. Wilkinson, Wilson & Ellemberg, 1997) since this would predict an effect of separation between edges of stimuli.

6.1.3 Opponents to the attentional hypothesis

On the other hand, some studies have not supported an attentional explanation of crowding. Wilkinson, Wilson & Ellemberg (1997) argue against an attentional hypothesis on the basis of finding no effect of pre-cueing a target on the effect of crowding during a contrast discrimination task. However, their targets were always presented at the centre of the display with flanking distracters, so it would not be surprising if observers simply directed their attention towards this object whether it was pre-cued or not.

More convincingly, Blake, Tadin, Sobel, Raissian and Chong (2006) argued for a role of primary visual cortical neurons in the crowding phenomenon. In contrast to He et al. (1996), they observed reductions in both the orientation-specific threshold elevation aftereffect and the motion aftereffect under crowding conditions. Most interestingly, they showed that the extent of adaptation reduction was correlated with the extent of the crowding effect. However, correlation cannot reveal the direction of causality, and these results are also of course consistent with crowding being caused by attentional processes beyond primary visual cortex which subsequently exert a top-down influence on the activity of cells in the primary visual cortex.

If crowding occurs due to limitations in low level feature representations, rather than attentional processes, it may be akin to lateral masking. Indeed, crowding and lateral masking have sometimes been used as interchangeable terms (e.g. Nazir, 1991). More recently, however, crowding as been conceptualised as a distinct form of lateral interaction (Chakravarthi & Cavanagh, 2007; Levi, Hariharan & Klein, 2002) despite possibly involving cells as early in the visual pathway as the primary visual cortex. For instance, Levi et al. (2002) demonstrated that feature detection (diagnostic of lateral masking) was still possible for stimuli where feature discrimination (diagnostic of crowding) was rendered very difficult by lateral flankers. Using the same detection-discrimination criterion, Chakravarthi et al. (2007) were able to show that crowding and lateral masking operate over different timescales.

Although it is clear that attention can relieve crowding (Felisberti, Solomon & Morgan, 2005; Huckauf & Heller, 2002; Morgan, Ward, Castet, 1998; Põder, 2005; Strasburger, 2005; Van der Lubbe & Keuss, 2001), there remains currently some debate as to whether this reflects a causal role for attention in crowding, or whether attention only acts to counteract the detrimental effects of crowding. The latter is a possibility since it is known that attention can enhance the precision with which the features of objects are represented in visual short-term memory (Wilken and Ma, 2004), and that it can enhance the signal in encoding of stimuli (Carrasco, Penpeci-Talgar & Eckstein, 2000; Lu & Dosher, 1998; Pestilli & Carrasco, 2005; Yeshurun & Carrasco, 1998). Signal enhancement may itself be accomplished by means of faster processing and therefore increased rates of information accrual (Carrasco & McElree, 2001). That attention may be able to counteract the effects of crowding, is also suggested by the fact that attention can act to reduce processing of task-irrelevant sources (e.g. Lu, Lesmes & Dosher, 2002; Pestilli & Carrasco, 2005), reflecting the 'limited capacity' notion of attention.

6.1.4 Attention, crowding and internal noise

Recently, Dakin, Bex, Cass, and Watt (2007) used equivalent noise methods to attempt to demonstrate that the effect of attention is qualitatively different from that of crowding in terms of noise reduction. In equivalent noise analysis, an attempt is made to partition out contributions from two sources of variability in performance: the sampling efficiency and the level of local noise in the system. When applied to vision science, the number of samples can be conceptualised as representing a modelled number of spatial or temporal samples taken by the visual system. Local noise values represent the contribution to variability made by the process of encoding stimuli. This includes any inefficiencies or interference in neural signalling or sources of imprecision as light is focussed in the eye, for instance. In the equivalent noise methodology, external noise is deliberately added to the stimulus by degrading the signal in some measurable manner. In the experiments of Dakin et al. (2007), observers reported the overall orientation of six log-Gabor filtered noise patches. Multiple orientations were superimposed at each of the six locations, and the variability of these orientations within a single location was manipulated to produce different levels of external noise. For each patch, each orientation was sampled from an underlying distribution with one of three standard deviations, the greater the standard deviation, the greater the level of external noise.

At low levels of external noise, we expect that the major source of variability in the system will be from "local" or "internal" noise. If we assume that external and internal noise act in an equivalent manner to affect performance, then as the level of external noise is increased, the relative contribution from internal noise will become less important. At high levels of external noise, differences in internal noise will have little effect. Equivalent noise analysis tells us:

$$\text{Thresholds} \propto \sqrt{\frac{\sigma^2_{\text{local}} + \sigma^2_{\text{external}}}{N_{\text{global}}}}$$

where N_{global} is the number of samples taken, σ_{local} is the standard deviation of local noise, and σ_{external} is the standard deviation of external noise. Example curves relating external noise and thresholds are given in Figure 25. They illustrate the effects of changing levels of internal noise and sampling on these curves. An increase in N corresponds to a downwards shift of the curve and a decrease in N corresponds to an upwards shift. A decrease in local noise corresponds to a decrease in thresholds at low external noise values, and less decrease in thresholds at high external noise values. An increase in local noise corresponds to large increases in thresholds at low external noise values and relatively less increase at high external noise values.

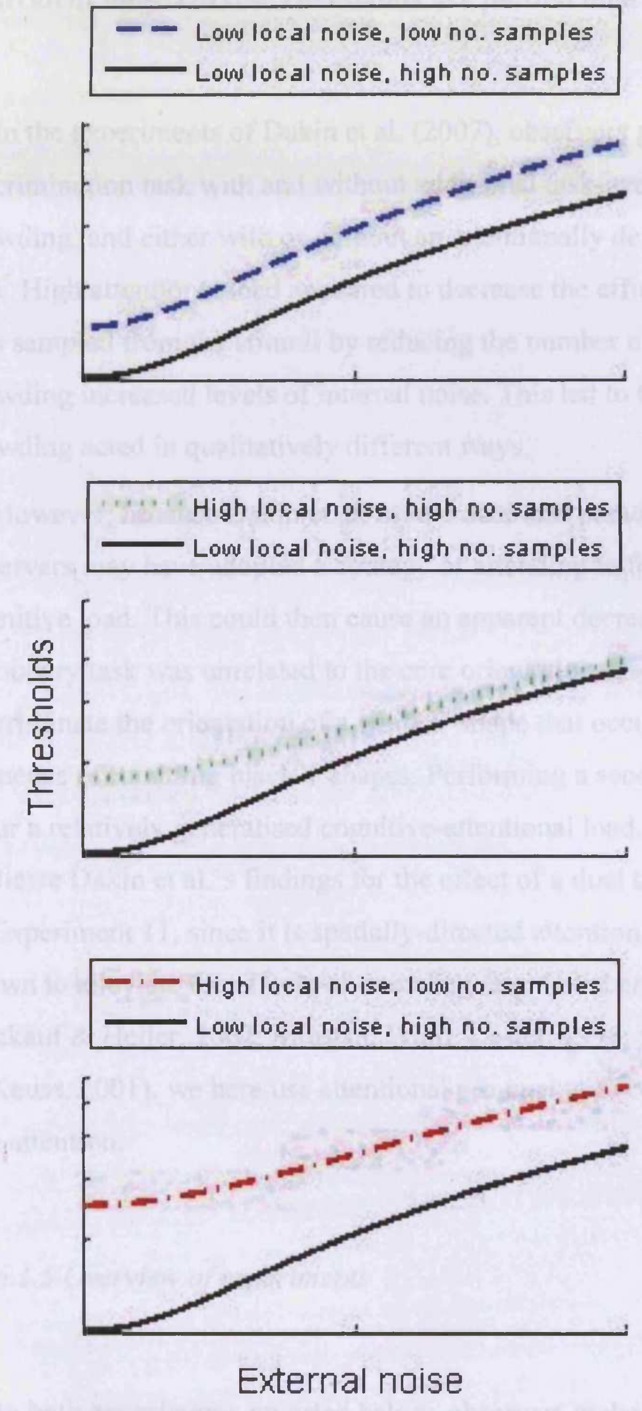
Low local noise, low no. samples
Low local noise, high no. samples
High local noise, high no. samples
Low local noise, high no. samples
High local noise, low no. samples
Low local noise, high no. samples
Threshods
External noise

Figure 25: the effect of changes in numbers of samples taken and of local noise on equivalent noise curves. Thresholds are plotted on a log ordinate axis.

In the experiments of Dakin et al. (2007), observers performed the orientation discrimination task with and without additional task-irrelevant distracters that produced crowding, and either with or without an attentionally demanding additional secondary task. High attentional load appeared to decrease the efficiency with which information was sampled from the stimuli by reducing the number of samples taken, whereas crowding increased levels of internal noise. This led to the conclusion that attention and crowding acted in qualitatively different ways.

However, because Dakin et al. used a dual task paradigm to increase attentional load, observers may have adopted a strategy of attending to fewer stimuli in order to reduce cognitive load. This could then cause an apparent decrease in sampling efficiency. Their secondary task was unrelated to the core orientation-discrimination task, and was to discriminate the orientation of a white T shape that occurred at a random time within a sequence of tumbling black T shapes. Performing a secondary task of this nature will incur a relatively generalised cognitive-attentional load. In Experiment 10, we attempt to replicate Dakin et al.'s findings for the effect of a dual task load on sampling efficiency. In Experiment 11, since it is spatially-directed attention that has most frequently been shown to alleviate the effects of crowding (e.g. Felisberti, Solomon & Morgan, 2005; Huckauf & Heller, 2002; Morgan, Ward, Castet, 1998; Strasburger, 2005; Van der Lubbe & Keuss, 2001), we here use attentional pre-cueing to compare the effects of crowding and attention.

6.1.5 Overview of experiments

In both experiments reported below, observers make global orientation discrimination judgements i.e. they report the 'overall' or 'average' orientation of several Gabor patches. In Experiment 10, we manipulate attention through the presence or absence of an additional secondary task requirement in a dual task method (high and low attentional load, respectively). In Experiment 11 the attentional manipulation is the presence or absence of pre-cues that specify which group of patches will later be queried (low and high attentional load, respectively). In addition, Experiment 11 also investigates the effect of crowding on the orientation discrimination judgement. Two preliminary experiments are also reported in Appendix 7.

6.2 EXPERIMENT 10 (DUAL TASK)

Experiment 10 was designed to measure the effect of a dual task attentional manipulation on orientation discrimination. We used similar primary and secondary tasks to those used by Dakin et al. (2007). Observers judged the mean or overall tilt of six Gabor patches (primary task) either in the presence or absence of a secondary task load. The secondary task was presented first in the sequence of trial events, but was secondary in the sense that it was introduced only in order to increase attentional load. To allow estimation of local internal noise levels, we additionally manipulated the variability of the orientation signal as a form of external noise.

6.3 METHOD FOR EXPERIMENT 10

A computer programme was written in Python using the VisionEgg library (http://www.visionegg.org) and displayed an array of stimuli against a mid-grey background on a 16 inch CRT screen refreshing at 75 Hertz. Observers viewed the display in a dimly lit room from a distance of 0.4m.

Throughout the experiment, observers fixated a small black square at the centre of the screen subtending 0.22 degrees. To ensure that observers were not unintentionally moving their eyes, eye position was constantly measured using a CRS video eyetracker, using the CRS Visage and videoEyetrace software. Trials where an eye movement of greater than approximately 0.5 degrees was detected were tagged as such in real-time by the experimenter.

6.3.1 Observers

Observers were five postgraduates at Cardiff University and one external participant, all of whom were naïve as to the purpose of the experiment, and four of whom were non-experts in participating in visual psychophysical experiments. All had normal or corrected-to-normal vision.

6.3.2 Procedure

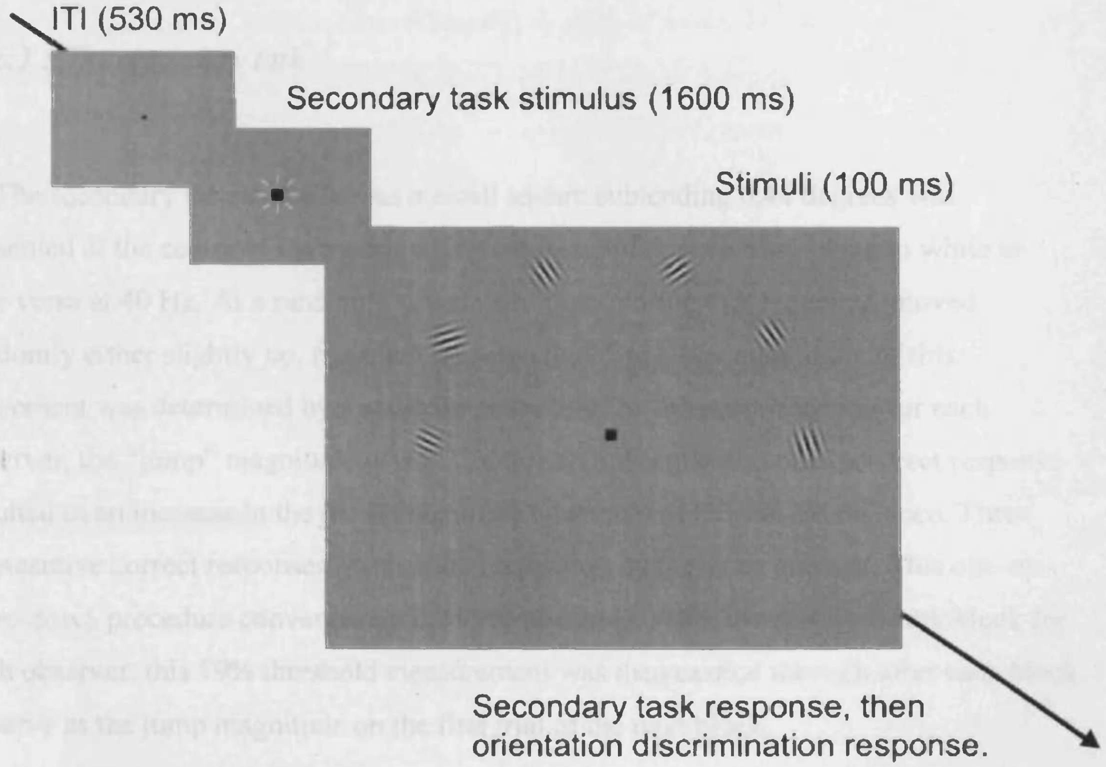

Figure 26: a typical trial timeline for Experiment 10. The secondary task stimulus is a flickering black-white square (40 Hertz). At a randomly determined time during its 1600 ms presentation, one cycle is presented at a displacement of x degrees either up, down, left or right of fixation, where the magnitude of x is determined by a staircase maintaining 79% correct.

At the start of the experiment, observers were given two practice blocks to familiarise themselves with the two different attentional conditions.

In general, we attempted to keep the details of the timing of stimuli close to those used by Dakin et al. (2007) to allow for as direct a comparison as possible. At the start of each trial, the fixation square was presented alone for 530 ms. Subsequently, depending on the attentional condition, either the fixation square or the secondary task stimulus was presented for a further 1600 ms.

6.3.3 The secondary task

The secondary task stimulus was a small square subtending 0.44 degrees was presented at the centre of the screen which changed luminance from black to white or vice versa at 40 Hz. At a randomly determined point during this sequence, moved randomly either slightly up, right, left or down for 25 ms. The magnitude of this movement was determined by a staircase procedure. At the start of testing for each observer, the "jump" magnitude was 0.22 degrees. Subsequently, one incorrect response resulted in an increase in the jump magnitude by a tenth of the current distance. Three consecutive correct responses resulted in a reduction by the same amount. This one-up three-down procedure converges on the 79% threshold. After the first dual task block for each observer, this 79% threshold measurement was then carried through after each block to serve as the jump magnitude on the first trial of the next block.

6.3.4 The primary task

The orientation discrimination stimuli were presented immediately after the offset of the secondary task stimulus. Six Gabor patches appeared, laid out in a semi-circular crescent such that two were immediately to the left and right, and others at intervals of 36 degrees and at a distance of 13.11 degrees from fixation. Gabors were presented for 100 ms. The presentation time of Gabors was short in order to minimise the likelihood of eye movements. The luminance of Gabors varied from 23.6 (trough) to 33.5 (peak) candelas per m^2. Gabors had a spatial period of 0.454 degrees per cycle and a variable orientation. The circular envelope that windowed the Gabor patches' amplitudes had a radius of 0.738 degrees of visual angle.

Tilt magnitudes for each Gabor were sampled from a Gaussian distribution with standard deviation either 0 10 or 20 degrees. By varying these standard deviations we were able to manipulate levels of external noise. The tilt directions represented by the mean of this Gaussian distribution (either left or right of vertical) were chosen randomly for each group on every trial. None of the Gabors was ever permitted a tilt magnitude greater than 45 degrees. The mean of the Gaussian distributions of orientations was determined by a staircase procedure. At the start of each block of trials, the mean of the tilt magnitude distributions was set at 20.0 degrees.

6.3.5 Response arrangements

Immediately after presentation of the Gabors, observers first responded either "up", "down", "left", or "right" for the secondary task stimulus using a keypress. In the low attentional load condition (no secondary task), observers simply pressed any of these keys. Immediately after this response, observers reported the global or average orientation of the six Gabors as either "right tilt from vertical" or "left tilt from vertical" using a keypress.

Two consecutive correct responses resulted in a reduction in the absolute mean of the Gaussian distributions of orientations (a reduction in overall tilts) by a tenth of the current absolute mean. One incorrect response caused an increase in the overall tilts by the same amount. This one-up two-down procedure converges on the 71% threshold.

6.3.6 Design

Within each block of trials, three separate orientation discrimination threshold measurements were taken with the staircase starting each time at five degrees. Each of the three measurements terminated after eight reversals in change of mean tilt magnitude. After three measurements had been taken, the block terminated.

Each block of the total 6 combinations of 2 x attention and 3 x external noise were each run in total three times resulting in 54 threshold measurements per observer. Within each of the three groups of 6 blocks, the order of conditions was randomised.

6.4 RESULTS FOR EXPERIMENT 10 (DUAL TASK)

Thresholds were determined for each staircase by taking the mean tilt magnitude over the last six reversals in tilt change direction (see Figure 27).

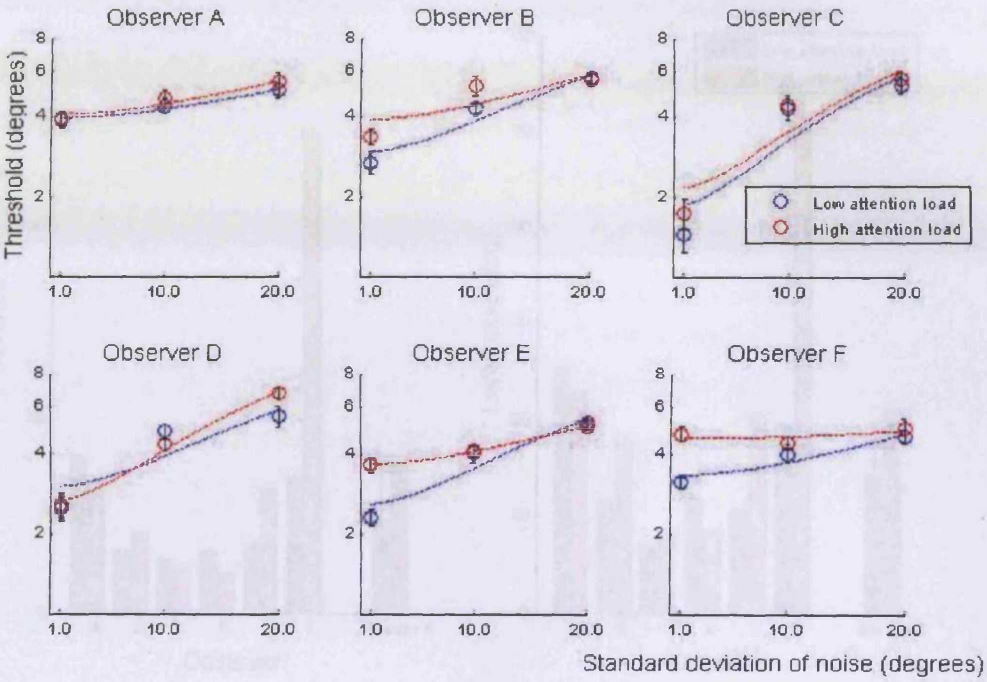

Figure 27: mean thresholds for each observer in each of the six total experimental conditions (3xexternal noise level and 2xattention). Error bars in black represent standard errors of the nine individual threshold measurements taken for each of the six conditions. Dotted lines represent curves fitted using the equivalent noise model.

There was a highly significant main effect of attention (F(1, 5) = 23.948, $p < 0.01$) and a significant effect of external noise level (F(2, 10) = 21.340, $p < 0.01$) in a 2 (attention level) x 3 (external noise level) repeated-measures ANOVA.

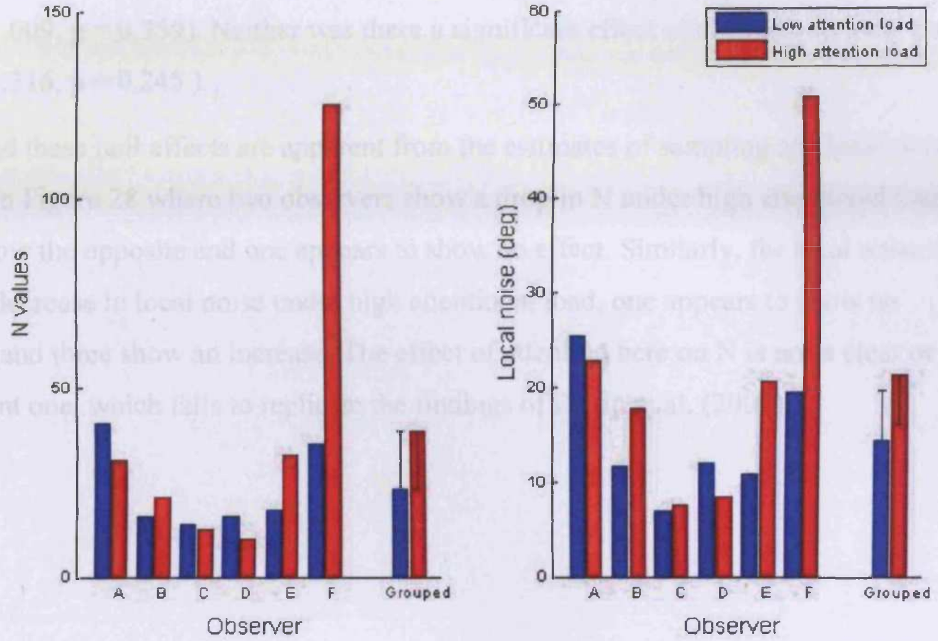

Figure 28: estimates of curve fitting to find N (sampling) and local noise values for each curve on the thresholds plot. Grouped values represent the means for all six observers and error bars here represent standard errors of within-subjects differences between the two conditions i.e. the difference between the conditions is calculated first for each observer, then the standard error of these differences is calculated. Error bars thus represent differences between conditions rather than variability within conditions.

From the thresholds shown in Figure 27, estimates of sampling and local noise values were derived using equivalent noise curve fitting (see Figure 28). The estimated values of the number of samples taken are frequently greater than the number of objects presented, suggesting that, if the assumptions of the equivalent noise equation are true, there are several samples per object. Each sample may represent one bar or trough in the Gabor patch, or the extent of a receptive field, or there may be several temporal samples during the viewing period.

There was no significant effect of attention on N (number of samples) in a paired t-test ($t(5) = 1.009$, $p = 0.359$). Neither was there a significant effect of attention on local noise ($t(5) = 1.316$, $p = 0.245$).

Indeed these null effects are apparent from the estimates of sampling and local noise plotted in Figure 28 where two observers show a drop in N under high attentional load, three show the opposite and one appears to show no effect. Similarly, for local noise two show a decrease in local noise under high attentional load, one appears to show no change, and three show an increase. The effect of attention here on N is not a clear or consistent one, which fails to replicate the findings of Dakin et al. (2007). [4]

---

[4] Two pilot experiments with very similar methods produced similar results to Experiment 1. In both, attention was manipulated using a similar dual task arrangement and orientation discrimination thresholds were measured. Estimates of sampling and local noise similarly showed no clear pattern. The methods and outcomes of these experiments are described in Appendix 7.

6.5 EXPERIMENT 11 (PRE-CUEING)

We failed in Experiment 10 to replicate the results of Dakin et al. (2007) for the effect of a dual task load on sampling. Experiment 11 had two purposes: to investigate whether a different type of attentional manipulation would yield a clearer pattern of change in the equivalent noise analysis, and to compare this with the effect of crowding. By introducing crowding as an additional variable, we were able to more directly compare the effects of attention and crowding when they are manipulated together. We were also able to test the second finding of Dakin et al. (2007) that crowding acts to increase local noise.

6.6 METHOD FOR EXPERIMENT 11

Attention was manipulated in Experiment 11 through spatial cueing instead of the use of a dual task. On half of trials, observers were given pre-cues indicating which of several groups of Gabors were targets (low attentional load). On the other half of trials there was uncertainty as to which of the five groups of Gabors was the target group on that trial (high attentional load). The methods for Experiment 11 were identical to those in Experiment 10 except for the following details.

6.6.1 Observers

Observers were two postgraduates at Cardiff University, two undergraduates at Cardiff University and two external participants. Five were naïve as to the purpose of the experiment, and four were non-experts in participating in visual psychophysical experiments. All had normal or corrected-to-normal vision.

6.6.2 Procedure

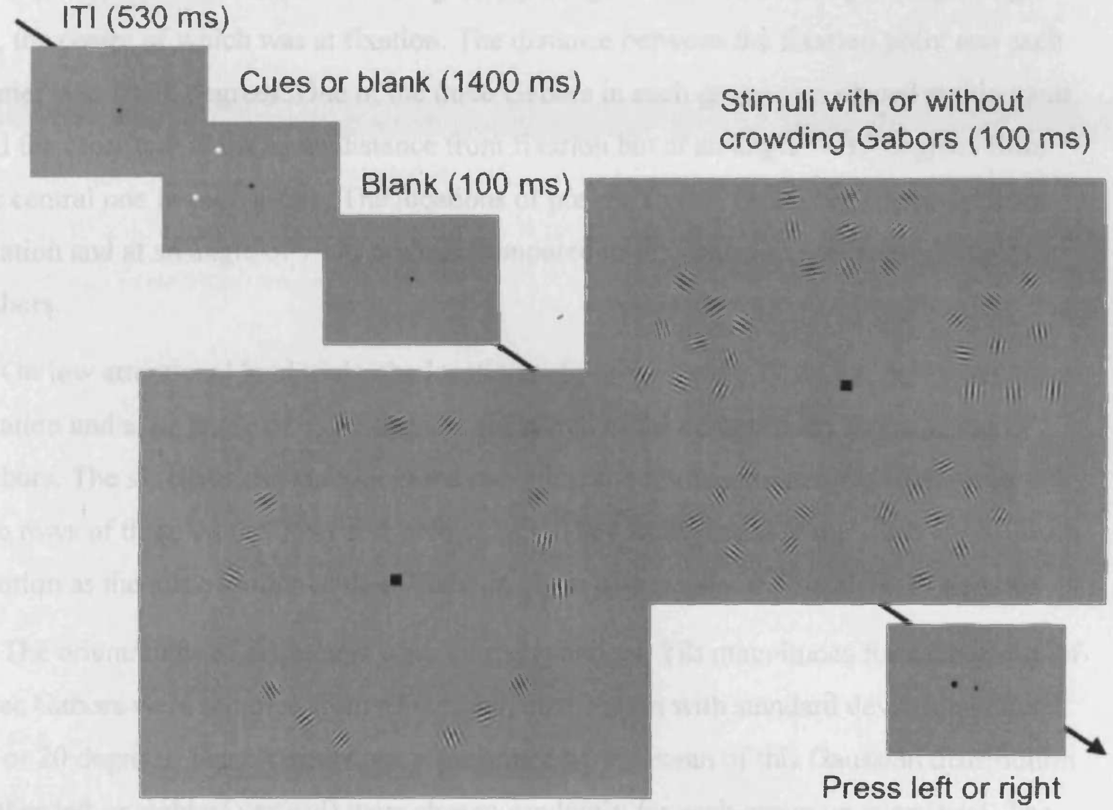

Figure 29: a typical trial timeline for Experiment 11.

At the start of the experiment, observers were given four practice blocks to familiarise themselves with the four crowding and attentional conditions.

At the start of each trial, the fixation square was presented alone for 530 ms and then the fixation square was presented for a further 1400 ms either alone or with a pair of attentional cues. The attentional cues were two white circles subtending a diameter of 0.73 degrees. Observers knew that the target Gabors would appear at the location in between these two markers. Following this, the fixation point was presented alone for 100 ms to prevent any possible forward masking of Gabors by the pre-cues.

Subsequently, five groups of three Gabor patches were presented either alone or with five groups of six additional distracter Gabors for 100 ms. The centres of the five groups were placed on five corners of an imaginary pentagon with one corner pointing straight up, the centre of which was at fixation. The distance between the fixation point and each corner was 10.98 degrees. One of the three Gabors in each group was placed at this point, and the other two at the same distance from fixation but at an angle +- 15 degrees from the central one in each group. The locations of pre-cues were 11.84 degrees away from fixation and at an angle of +- 25 degrees compared to the centre of the target group of Gabors.

On low attentional load trials, the locations of pre-cues were 11.84 degrees away from fixation and at an angle of +- 25 degrees compared to the centre of the target group of Gabors. The six distracter Gabors in the crowding condition consisted for each group of two rows of three on the inner and outer edges. They were placed at the same angles from fixation as the three groups of three Gabors, but at distances of 8.56 and 13.37 degrees.

The orientations of distracters were entirely random. Tilt magnitudes for each group of three Gabors were sampled from a Gaussian distribution with standard deviation either 0 10 or 20 degrees. The tilt directions represented by the mean of this Gaussian distribution (either left or right of vertical) were chosen randomly for each group on every trial. The mean of the Gaussian distributions of orientations was determined by a staircase procedure. At the start of each block of trials, the mean of the tilt magnitude distributions for Gabors in all five groups was set at 20.0 degrees.

At the end of every trial, the post-cue immediately indicated one group of Gabors to be reported. The post-cue was a black circle of radius 0.47 degrees and was presented 4.44 degrees eccentric from fixation and in the same direction from fixation as the centre of the group of Gabors. Of course, in the low attentional load condition, observers already knew which group they would report and the pre-cues rendered this post-cue redundant. The target group of Gabors was determined randomly on every trial. Observers then reported the global orientation of the three post-cued Gabors.

6.6.3 Design

We used a 2 x 2 within-subjects design: crowding versus no crowding and attentional cues versus no attentional cues. For each of these four combinations there were in addition three possible levels of external noise. Each of the total 12 combinations of 2 x attention, 2 x crowding and 3 x external noise were each run in total three times resulting in 36 threshold measurements per observer. Within each of the three groups of 12 blocks, the order of conditions was randomised.

6.7 RESULTS FOR EXPERIMENT 11 (PRE-CUEING)

Thresholds were determined for each block by taking the mean tilt magnitude over the last ten reversals in tilt change direction.

○ Low attention load, no crowding
○ High attention load, no crowding
○ Low attention load, crowding
○ High attention load, crowding

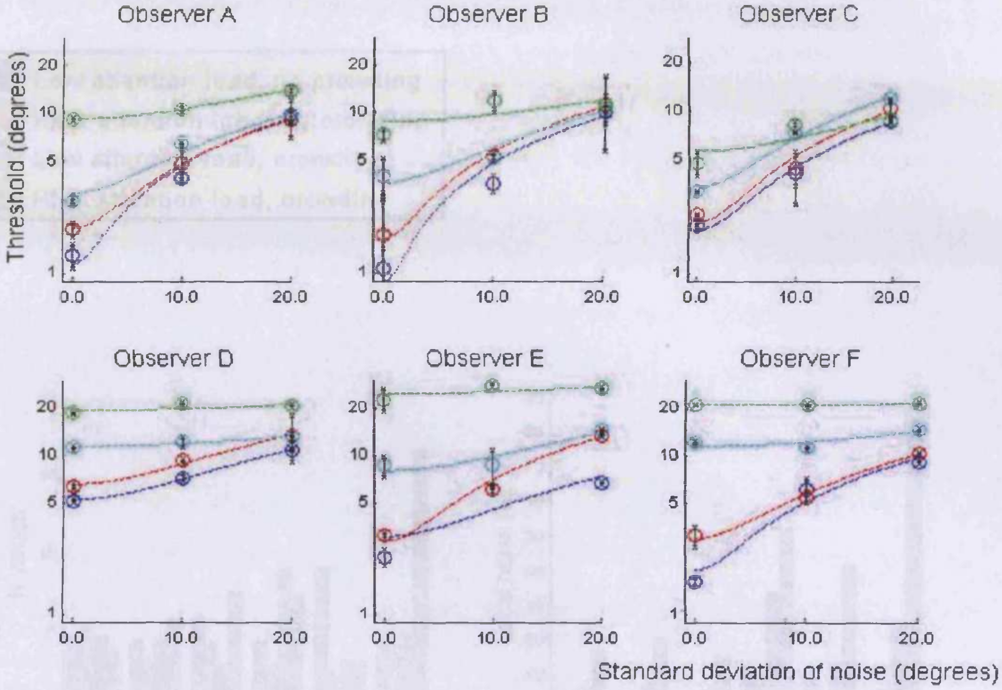

Figure 30: mean thresholds for each observer in each of the twelve total experimental conditions (3xexternal noise, 2xattention and 2xcrowding). Error bars in black represent standard errors obtained by bootstrapping. Dotted lines represent curves fitted using the equivalent noise model.

There were significant main effects of all three experimental variables in a 2 (attention level) x 2 (crowding level) x 3 (external noise level) repeated-measures ANOVA (effect of attention: $F(1, 5) = 10.904$, $p < 0.05$ effect of crowding: $F(1, 5) = 14.107$, $p < 0.05$ effect of external noise: $F(2, 10) = 114.806$, $p < 0.01$).

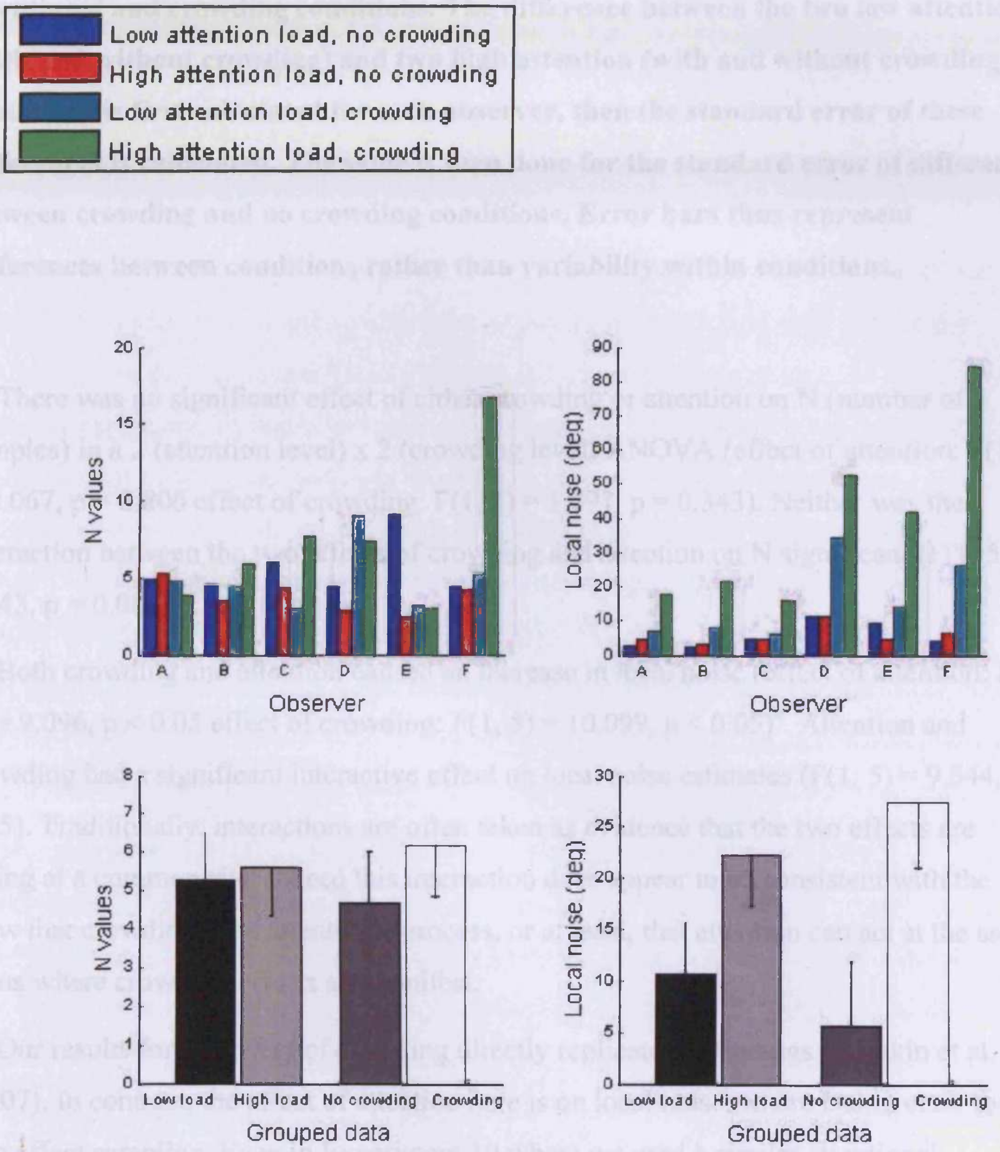

Figure 31: estimates from the results of curve fitting to find N (sampling) and local noise values for each curve on the thresholds plot. Grouped values represent the means for all six observers for all low attention trials and all high attention trials (irrespective of crowding), all trials without crowding and all trials with crowding (irrespective of attentional load), respectively i.e. they show the main effects. Error

bars here represent standard errors of within-subjects differences between the two attentional and crowding conditions. The difference between the two low attention (with and without crowding) and two high attention (with and without crowding) conditions is first calculated for each observer, then the standard error of these differences is calculated. The same is then done for the standard error of differences between crowding and no crowding conditions. Error bars thus represent differences between conditions rather than variability within conditions.

There was no significant effect of either crowding or attention on N (number of samples) in a 2 (attention level) x 2 (crowding level) ANOVA (effect of attention: $F(1, 5) = 0.067$, $p = 0.806$ effect of crowding: $F(1, 5) = 1.097$, $p = 0.343$). Neither was the interaction between the two effects of crowding and attention on N significant ($F(1, 5) = 4.643$, $p = 0.084$).

Both crowding and attention caused an increase in local noise (effect of attention: $F(1, 5) = 9.096$, $p < 0.05$ effect of crowding: $F(1, 5) = 10.099$, $p < 0.05$)[5]. Attention and crowding had a significant interactive effect on local noise estimates ($F(1, 5) = 9.544$, $p < 0.05$). Traditionally, interactions are often taken as evidence that the two effects are acting at a common site. Indeed this interaction does appear to be consistent with the view that crowding is an attentional process, or at least, that attention can act at the same locus where crowding effects are manifest.

Our results for the effect of crowding directly replicate the findings of Dakin et al. (2007). In contrast, the effect of attention here is on local noise, where Dakin et al. found it to affect sampling. Even in Experiment 10 where we used a similar attentional manipulation to Dakin et al., we found no effect on sampling.

---

[5] One plausible constraint that could be imposed during the curve fitting process is that thresholds should rise monotonically with levels of external noise. Curve fitting outcomes were only very slightly different for both experiments when this constraint was imposed. The methods used and outcomes of curve fitting are described in Appendix 8.

6.8 DISCUSSION OF EXPERIMENTS 10 AND 11

In Experiment 10 we failed to replicate the results found by Dakin et al. (2007) i.e. that the effect of the attentional load imposed by a secondary task was to decrease the sampling of stimuli. Instead, Experiment 10 showed no coherent effect of a dual task manipulation on either sampling efficiency or local noise. We find this, despite the secondary task having a significantly detrimental effect on performance. In addition, we find no clear pattern in two pilot experiments described in Appendix 7 of the effects of a dual task manipulation on sampling or local noise. In contrast, in Experiment 11 we found that both spatial cueing and crowding affected local noise values, with crowding acting to increase local noise and spatial attention acting to decrease local noise. Thus it appears that the type of attentional manipulation was critical in Experiments 10 and 11, with spatial attention having a more specific type of effect than general cognitive load. In addition, in contrast to Dakin et al., the qualitative effects of crowding and attention were not dissociable in terms of local noise and sampling when spatial cueing was used.

In our Experiment 10 we found a significant effect of attention on orientation discrimination when attention was manipulated through the presence or absence of a secondary task. This effect however, when analysed in terms of equivalent noise, did not turn out to be specific either to attentional sampling nor to levels of internal noise. Our task was very similar to that used by Joseph, Chun and Nakayama (1997) who examined the extent to which orientation is a ‘preattentive’ feature. They manipulated attention using the secondary task of reporting the identity of a single white letter amid a stream of rapidly presented black letters. They reasoned that if, in fact, features like colour and orientation are indeed ‘preattentive’ and thus accessible even in the absence of directed attention, then this secondary task should have no effect on detection of an orientation singleton amongst homogenously oriented distracters. In fact they did find an effect of the secondary task on performance, thus contradicting the idea that orientation is always ‘preattentively’ accessible. Our null effect of the secondary task on sampling and internal noise in Experiment 10 suggests that the effects found by Joseph et al. may have been of a similar coarse and unspecific nature i.e. that this kind of attentional manipulation does not offer a 'pure' effect on orientation discrimination.

6.8.1 Effects of attention on sampling

In the experiments of Dakin et al. (2007), the effect of the dual task manipulation was to cause a decrease in the number of samples taken. This could, perhaps, have been due to a) a reduction in the total available attentional resource for sampling, or b) a reduction in the capacity of attention to move serially between stimuli. The former (a) is a possibility since attention suffers from being shared between many objects: as shown in Chapter 2, performance decrements in reporting orientations were observed when the number of attended objects was increased from one to four, and even from one to two. The latter (b) is a possibility since when attention is required to move serially, it takes in the order of 100-200 ms to move between objects (Wolfe, Alvarez & Horowitz, 2000). Gabors were presented immediately after the secondary task stimulus and for only 100 ms. Attending first to a secondary task might mean that the Gabors disappeared during the time attention was moving towards or between them, thus reducing the number of samples taken. In Experiment 10 here, however, it seems that the dual task did not offer a 'pure' effect, perhaps due to encoding and retrieval components associated with the task. For instance, our secondary task is likely to have involved encoding the direction of motion of the flickering square into memory, holding it in memory through the presentation of the primary task stimuli, and then retrieving the direction of motion for report while holding the primary task response in memory, All of these elements are likely to have affected performance on the primary task in some way. That we find a clearer pattern for cueing than a dual task is consistent with the fact that it is spatially-directed attention that has most frequently been shown to alleviate the effects of crowding (e.g. Felisberti, Solomon & Morgan, 2005; Huckauf & Heller, 2002; Morgan, Ward, Castet, 1998; Strasburger, 2005; Van der Lubbe & Keuss, 2001).

6.8.2 Effects of attention on internal noise

In Experiment 11, according to the equivalent noise analysis, spatial cueing has its effect through raising the precision with which orientations are represented, or equivalently by lowering internal noise. This is consistent with the findings of Lu and Dosher (1998) for an orientation discrimination task. They varied the contrast of random external noise to assess the effect of cueing on local noise. The observed that attention had a greater effect when external noise was low but that it had minimal effect when external noise was high, a pattern characteristic of local internal noise effects. Here we are able to replicate their result, and through the use of equivalent noise methods we are also able to make quantitative estimates of internal noise values. It is also consistent with the findings presented in Chapter 2, that pre-cueing to one of a number of potential targets increases the precision with which either the orientations, spatial periods or positions of objects are represented, when observers are required to continuously monitor these changing features. In Experiment 11, the noise reduction benefit of the spatial cue may have acted to reduce noise during encoding of stimuli at the target locations or noise from other locations on the screen, the latter of which has been shown recently to make up a large proportion of attentional cuing benefits (Gould, Wolfgang & Smith, 2007).

Although we did not observe an effect of a dual task manipulation on internal noise, Prinzmetal, Amiri, Allen and Edwards (1998) showed in a series of experiments that attention reduced the noisiness in encoding of responses in a task where observers reported the features of briefly presented objects, either in the presence or absence of a secondary task. It is not clear why they found an effect of attention on internal noise using a dual task design, where we did not in Experiment 10 and in the two experiments described in Appendix 7. It seems possible that this may be to do with the nature and demands of the secondary task used: theirs involved a serial search amid a 3x3 array of matrix of letters to identify either a T or an F within the array. On average then, observers would have to search four or five objects in the array. Given that the displays were presented for a period in the order of 30-100 ms, and that visual search is known to proceed at rates in the order of 40 ms per item (e.g. Wolfe, Alvarez & Horowitz, 2000), this is clearly a demanding task. It may have meant that attention did not have time to move between the secondary task and the target stimulus before offset of the stimuli. This may have meant that their target stimuli were hardly attended at all on some trials, causing this increase in the spread of responses under dual task conditions. That attention is associated with noise reduction, was clear, however.

6.8.3 Relation to effects of attention on visual short-term memory

Finally, our result in Experiment 11 is also consistent with similar findings of Wilken and Ma (2004) for visual short-term memory (VSTM) tasks. In the experiments of Wilken and Ma, observers were presented with a variable number of colour, orientation or spatial frequency stimuli. After offset of the display and 1500 ms retention interval, they were prompted to report the appropriate feature of one of the remembered objects. They found that the spread of errors made in this report depended on the number of objects attended i.e. that the effect of attention was to increase precision of the representation in memory, or equivalently, to reduce the noise in this representation. That the effect of attention on short-term memory is similar to its effect on perception is of interest as the capacity limits for VSTM and visual attention may be intimately related (e.g. Cowan, 2000).

6.9 SUMMARY OF FINDINGS IN CHAPTER 6

Contrary to the findings of Dakin et al. (2007), we do not find qualitatively different types of effects of attention and crowding, instead, both here have their effects through local noise and not sampling. Whilst it is still a logical possibility that crowding and attention are separable processes that both simply happen to exhibit effects on local noise, our results here rule out their separation on the grounds of effects on sampling and local noise. Indeed these findings are not incompatible with the view that crowding occurs when attention fails to select a target to the exclusion of distracters (e.g. He, Cavanagh &

Intriligator, 1996; Intriligator & Cavanagh, 2001; Tripathy & Cavanagh, 2002). Future research might, for instance, attempt to use these equivalent noise methods to cast further light on the issue of whether low-level processes such as adaptation are affected by crowding (Blake, Tadin, Sobel, Raissian & Chong, 2006), or whether crowding occurs at a later stage more consistent with an attentional selection stage (He, Cavanagh & Intriligator, 1996). The estimation of two separate parameters (sampling and local noise) instead of a single performance measure may provide a more sensitive test of the nature of the effects observed.

7. SUMMARY AND FINAL REMARKS

In this thesis, two main types of paradigm have been employed. In Chapters 2 – 5, a continuous monitoring paradigm was used to investigate the nature of the capacity limit for attending to the positional and non-positional features of multiple objects. Three main conclusions can be drawn from this work.

Firstly, that performance in Chapter 2 was not dependent on a previously suggested limit based on a fixed number of objects, where this number of objects is around four. Rather, precision of representations of the features of objects exhibited a decline before this limit, even between tracking one object and tracking two objects in the case of orientation and position tracking.

Secondly, tracking responses exhibited perceptual lags. That is, responses were more similar to values of the queried object at a time before offset, rather than the actual last value displayed. In Chapter 2, for orientation, and particularly for spatial period, perceptual lags increased as more objects were added to the attentional load. Analyses of error patterns broken down by likelihoods of drawing from a guessing strategy, suggested a contribution from serial processing. In Chapter 3, two double report experiments provided further support for the presence of a serial element to processing of spatial periods.

Thirdly, different features exhibited different patterns of lags. Greater lags and more effect of set size on lags were observed for spatial period than orientation or position tracking. Three pieces of evidence all point to the serial processing component being greatest when monitoring spatial periods. The lags and set size effect on lags were greatest by far in the case of spatial period (Chapter 2), the correlation between errors was smallest (Chapter 3), and the analysis of responses by likelihood of drawing from a guessing strategy, also yielded the largest difference (Chapter 2). Four experiments in Chapter 4 ruled out the possibilities that the differences in lags between orientation and spatial period were due to artefacts relating to the range of feature values. Possible reasons why spatial period tracking was different from orientation and position tracking were discussed in Chapter 5. We considered whether observers in fact used a position tracking strategy to perform orientation tracking, and whether this might mean that orientation and position tracking were performed by the dorsal processing stream in the brain, where spatial period was performed by the ventral stream. This remains an open question, however, and deserves further investigation.

In Chapter 6, a different methodology was used to investigate the benefits of attentional cueing in terms of noise reduction and sampling. Pre-cueing and dual task manipulations were used to compare the effects of attention with those of crowding, which has been demonstrated to increase internal noise. Contrary to previous studies, the effect of a secondary task was not found to be specific to a reduction in attentional sampling. Rather, in Experiment 10 and in two preliminary experiments described in Appendix 7, a secondary task raised orientation discrimination thresholds without this effect being specific to either sampling reduction or internal noise inflation. Spatially directed attention, however, had an overall effect of lowering thresholds, and was found specifically to reduce internal noise, consistent with the precision results of the previous chapters. In this respect, since crowding has been previously shown to increase internal noise, an effect that we replicated in Experiment 11, crowding and spatially directed attention were not separable here in terms of the locus of their action.

In the continuous monitoring tasks presented in Chapters 2 - 4, we showed that attention increased the precision with which the features of objects were represented. One advantage of reporting stimulus features directly, is that differences in the distribution of errors should reflect differences in the precision of encoding for different numbers of objects attended. This measure is uncontaminated by the kind of decision noise discussed in Chapter 2 where noise that arises when individual responses are based on the representations of more than one object. Future studies could implement an equivalent noise paradigm for a continuous monitoring task with attention to the features of multiple objects. Equivalent noise methods, when applied to the continuous monitoring paradigm, may reveal that attentional effects comprise an effect on sampling, which may be contributing to the observed precision results. In turn, this sampling, if temporal in nature, could provide further evidence for a contribution from serial processing.

REFERENCES

Alvarez, G. A. & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science, 15*(2), 106-111.

Alvarez, G. A. & Cavanagh, P. (2005). Independent resources for attentional tracking in the left and right visual hemifields. *Psychological Science, 16*(8), 637-643.

Alvarez, G. A. & Franconeri, S. L. (2007). How many objects can you track? Evidence for a resource-limited attentive tracking mechanism. *Journal of Vision, 7*(13), 1-10.

Alvarez, G. A., Horowitz, T. S., Arsenio, H. C., DiMase, J. S. & Wolfe, J. M. (2005). Do multielement visual tracking and visual search draw continuously on the same visual attention resources? *Journal of Experimental Psychology: Human Perception and Performance, 31*(4), 643-667.

Awh, E., Barton, B. & Vogel, E. K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science, 18*(7), 622-628.

Bahrami, B. (2003). Object property encoding and change blindness in multiple object tracking. *Visual Cognition, 10*(8), 949-963.

Bichot, N. P., Cave, K. R. & Pashler, H. (1999). Visual selection mediated by location: feature-based selection of noncontiguous locations. *Perception and Psychophysics, 61*(3), 403-423.

Blackmore, S. J., Brelstaff, G., Nelson, K. & Troscianko, T. (1995). Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception, 24*(9), 1075-1081.

Blake, R., Tadin, D., Sobel, K. V., Raissian, T. A. & Chong, S. C. (2006). Strength of early visual adaptation depends on visual awareness. *Proceedings of the National Academy of Sciences of the United States of America, 103*, 4783-4788.

Blakemore, C., Nachmias, J., Sutton, P. (1970). The perceived spatial frequency shift: evidence for frequency-selective neurones in the human brain. *Journal of Physiology, 210*, 727-750.

Blaser, E., Pylyshyn, Z. W. & Holcombe, A. O. (2000). Tracking an object through feature space. *Nature, 408*(6809), 196-199.

Bouma, H. (1970). Interaction effects in parafoveal letter recognition. *Nature, 226*, 177-178.

Bouma, H. (1973). Visual interference in the parafoveal recognition of initial and final letters of words. *Vision Research, 13*, 767-782.

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

Broadbent, D. (1958). *Perception and communication.* London: Pergamon Press.

Broadbent, D. E., Gregory, M. (1963). Division of attention and the decision theory of signal detection. *Proceedings of the Royal Society of London B, 158*(971), 222-231.

Carlson, T. A., Hogendoorn, H. & Verstraten, F. A. J. (2006). The speed of visual attention: what time is it? *Journal of Vision, 6*, 1406-1411.

Carlson, T., VanRullen, R., Hogendoorn, H., Verstraten, F. & Cavanagh, P. (2007). Distinguishing models of multifocal attention: it's a matter of time [abstract]. *Journal of Vision, 7*(9), 641a.

Carrasco, M. & McElree, B. (2001). Covert attention accelerates the rate of visual information processing. *Proceedings of the National Academy of Sciences of the United States of America, 98*(9), 5363-5367.

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

Chakravarthi, R. & Cavanagh, P. (2007). Temporal properties of the polarity advantage effect in crowding. *Journal of Vision, 7*(2), 1-13.

Cheal, M. & Lyon, D. R. (1991). Central and peripheral precuing of forced-choice discrimination. *The Quarterly Journal of Experimental Psychology Section A: Human Experimental Psychology, 43*(4), 859-880.

Cherry, E. C. (1953). Some experiments on the recognition of speech, with one and with two ears. *The Journal of the Acoustical Society of America, 25*(5), 975-979.

Cowan, N. (2000). The magical number 4 in short-term memory: a reconisderation of mental storage capacity. *Behavioral and Brain Sciences, 24*, 87-185.

Culham, J. C., Brandt, S. A., Cavanagh, P., Kanwisher, N. G., Dale, A. M. & Tootell, R. B. H. (1998). Cortical fMRI activation produced by attentive tracking of moving targets. *Journal of Neurophysiology, 80*(5), 2657-2670.

Dakin, S. C., Bex, P. J., Cass, J. & Watt, R. J. (2007). No role for attention in orientation crowding [abstract]. *Perception, 36*, 303-304.

Deutsch, J. A. & Deutsch, D. (1963). Attention: some theoretical considerations. *Psychological Review, 70*(1), 51-61.

Duncan, J. & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96*(433-458).

Duncan, J. & Humphreys, G. (1992). Beyond the search surface: visual search and attentional engagement. *Journal of Experimental Psychology: Human Perception and Performance, 18*(2), 578-588.

Duncan, J., Ward, R. & Shapiro, K. (1994). Direct measurement of attentional dwell time in human vision. *Nature, 369*, 313-315.

Felisberti, F. M., Solomon, J. A. & Morgan, M. J. (2005). The role of target salience in crowding. *Perception, 34*, 823-833.

Franconeri, S. L., Alvarez, G. A. & Enns, J. T. (2007). How many locations can be selected at once? *Journal of Experimental Psychology: Human Perception and Performance, 33*, 1003-1012.

Freyd, J. J. & Finke, R. A. (1984). Representational Momentum. *Journal of Experimental Psychology: Learning, Memory and Cognition, 10*(1), 126-132.

Gould, I. C., Wolfgang, B. J. & Smith, P. L. (2007). Spatial uncertainty explains exogenous and endogenous attentional cuing effects in visual signal detection. *Journal of Vision, 7*(13), 1-17.

He, S., Cavanagh, P. & Intriligator, J. (1996). Attentional resolution and the locus of visual awareness. *Nature, 383*, 334-337.

Horowitz, T. S., Holcombe, A. O., Wolfe, J. M., Arsenio, H. C. & DiMase, J. S. (2004). Attentional pursuit is faster than attentional saccade. *Journal of Vision, 4*, 585-603.

Horowitz, T. S., Place, S. S., Van Wert, M. J. & Fencsik, D. E. (2007). The nature of capacity limits in multiple object tracking [abstract]. *Perception, 36*, 9.

Huang, L. & Pashler, H. (2007). A Boolean map theory of visual attention. *Psychological Review, 114*(3), 599-631.

Huang, L., Treisman, A. & Pashler, H. (2007). Characterizing the limits of human visual awareness. *Science, 317*, 823-825.

Huckauf, A. & Heller, D. (2002). Spatial selection in peripheral letter recognition: in search of boundary conditions. *Acta Psychologica, 111*, 101-123.

Intraub, H. (1997). The representation of visual scenes. *Trends in Cognitive Sciences, 1*(6), 217-221.

Intriligator, J. & Cavanagh, P. (2001). The spatial resolution of visual attention. *Cognitive Psychology, 43*, 171-216.

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

Joseph, J. S., Chun, M. M., Nakayama, K. (1997). Attentional requirements in a 'preattentive' feature search task. *Nature, 387*, 805-807.

Jovicich, J., Peters, R. J., Koch, C., Braun, J., Chang, L. & Ernst, T. (2001). Brain areas specific for attentional load in a motion-tracking task. *Journal of Cognitive Neuroscience, 13*(8), 1048-1058.

Keane, B. P. & Pylyshyn, Z. W. (2006). Is motion extrapolation employed in multiple object tracking? Tracking as a low-level, non-predictive function. *Cognitive Psychology, 52*, 346-368.

Kerzel, D. (2000). Eye movements and visible persistence explain the mislocalization of the final position of a moving target. *Vision Research, 40*, 3703-3715.

Kerzel, D. (2006). Why eye movements and perceptual factors have to be controlled in studies on "representational momentum". *Psychonomic Bulletin and Review, 13*(1), 166-173.

Khurana, B., Watanabe, K. & Nijhawan, R. (2000). The role of attention in motion extrapolation: are moving objects 'corrected' or flashed objects attentionally delayed? *Perception, 29*, 675-692.

Levene, H. (1960). Robust tests for equality of variances. In I. Olkin (Ed.), *Contributions to probability and statistics: essays in honor of Harold Hotelling* (pp. 278-292). Palo Alto, California: Stanford University Press.

Levi, D. M., Hariharan, S. & Klein, S. A. (2002). Suppressive and facilitatory spatial interactions in peripheral vision: peripheral crowding is neither size invariant nor simple contrast masking. *Journal of Vision, 2*, 167-177.

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

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

Lu, Z.-L., Lesmes, L. A. & Dosher, B. A. (2002). Spatial attention excludes external noise at the target location. *Journal of Vision, 2*(4), 312-323.

Lu, Z., Neuse, J., Madigan, S. & Dosher, B. A. (2005). Fast decay of iconic memory in observers with mild cognitive impairments. *Proceedings of the National Academy of Sciences of the United States of America, 102*(5), 1797-1802.

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

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

Magnussen, S. & Greenlee, M. W. (1997). Competition and sharing of processing resources in visual discrimination. *Journal of Experimental Psychology: Human Perception & Performance, 23*(6), 1603-1616.

Mäkelä, P., Whitaker, D. & Rovamo, J. (1993). Modelling of orientation discrimination across the visual field. *Vision Research, 33*(5/6), 723-730.

McElree, B. & Carrasco, M. (1999). The temporal dynamics of visual search: evidence for parallel processing in feature and conjunction searches. *Journal of Experimental Psychology: Human Perception and Performance, 25*(6), 1517-1539.

Milner, A. D. & Goodale, M. A. (1995). *The visual brain in action.* Oxford: Oxford University Press.

Mishkin, M., Ungerleider, L. G. & Macko, K. A. (1983). Object vision and spatial vision: two cortical pathways. *Trends in Neuroscience, 6*, 414-417.

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, 51*(2), 347-370.

Movshon, J. A., Thompson, I. D., Tolhurst, D. J. (1978). Receptive field organization of complex cells in the cat's striate cortex. *Journal of Physiology, 283*, 79-99.

Nakayama, K. & Mackeben, M. (1989). Sustained and transient components of focal visual attention. *Vision Research, 29*(11), 1631-1647.

Nazir, T. A. (1992). Effects of lateral masking and spatial precueing on gap-resolution in central and peripheral vision. *Vision Research, 32*(4), 771-777.

Nijhawan, R. (1994). Motion extrapolation in catching. *Nature, 370*, 256-257.

Nissen, M. J., Corkin, S. (1985). Effectiveness of attentional cueing in older and younger adults. *Journal of Gerontology, 40*, 185-191.

Ogmen, H., Breitmeyer, B. G., Melvin, R. (2003). The what and where in visual marking. *Vision Research, 43*, 1337-1350.

O'Regan, J. K. (1992). Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Experimental Psychology, 46*(3), 461-488.

O'Regan, J. K., Rensink, R. A. & Clark, J. J. (1999). Change blindness as a result of "mudsplashes". *Nature, 398*(6722), 34.

Palmer, J. (1995). Attention in visual search: distinguishing four causes of a set-size effect. *Current Directions in Psychological Science, 4*(4), 118-123.

Pelli, D. G., Palomares, M. & Majaj, N. J. (2004). Crowding is unlike ordinary masking: Distinguishing feature integration from detection. *Journal of Vision, 4*, 1136-1169.

Perenin, M.-T. & Vighetto, A. (1988). Optic ataxia: a specfic disruption in visuomotor mechanisms. I. Different aspects of the deficit in reaching for objects. *Brain, 111*, 643-674.

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

Põder, E. (2005). Effect of attention on the detection and identification of masked spatial patterns. *Perception, 34*, 305-318.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32*, 3-25.

Posner, M. I., Petersen, S. E. (1990). The attention sytem of the human brain. *Annual Review of Neuroscience, 13*, 25-42.

Prinzmetal, W., Amiri, H., Allen, K. & Edwards, T. (1998). Phenomenology of attention: 1. color, location, orientation and spatial frequency. *Journal of Experimental Psychology: Human Perception and Performance, 24*(1), 261-282.

Pylyshyn, Z. W. (1989). The role of location indexes in spatial perception: a sketch of the FINST spatial-index model. *Cognition, 32*, 65-97.

Pylyshyn, Z. W., Storm, R. W. (1988). Tracking multiple independent targets: evidence for a parallel tracking mechanism. *Spatial Vision, 3*(3), 1-19.

Reeves, A. & Sperling, G. (1986). Attention gating in short-term visual memory. *Psychological Review, 93*(2), 180-206.

Regan, D. (2000). *Human perception of objects*: Sinauer Publishers.

Rensink, R. A., O'Regan, J. K. & Clark, J. J. (1997). To see or not to see: the need for attention to perceive changes in scenes. *Psychological Science, 8*, 368-373.

Richter, E. S. & Yager, D. (1984). Spatial-frequency difference thresholds for central and peripheral viewing. *Journal of the Optical Society of America A, 1*(11), 1136-1139.

Roulston, B. W., Self, M. W. & Zeki, S. (2006). Perceptual compression of space through position integration. *Proceedings of the Royal Society of London B, 273*(1600), 2507-2512.

Saiki, J. (2003a). Feature binding in object-file representations of multiple moving items. *Journal of Vision, 3*, 6-21.

Saiki, J. (2003b). Spatiotemporal characteristics of dynamic feature binding in visual working memory. *Vision Research, 43*, 2107-2123.

Schiller, P. H. (1966). Forward and backward masking as a function of relative overlap and intensity of test and masking stimuli. *Perception and Psychophysics, 1*(5), 161-164.

Scholl, B. J. (2001). Objects and attention: the state of the art. *Cognition, 80*, 1-46.

Scholl, B. J. & Pylyshyn, Z. W. (1999). Tracking multiple items through occlusion: clues to visual objecthood. *Cognitive Psychology, 38*, 259-290.

Scholl, B. J., Pylyshyn, Z. W. & Feldman, J. (2001). What is a visual object? Evidence from target merging in multiple object tracking. *Cognition, 80*, 159-177.

Scholl, B. J., Pylyshyn, Z. W. & Franconeri, S. L. (1999). When are spatiotemporal and featural properties encoded as a result of attentional allocation? [abstract]. *Investigative Ophthalmology and Visual Science, 40*(4), S797.

Sheth, B. R., Nijhawan, R. & Shimojo, S. (2000). Changing objects lead briefly flashed ones. *Nature Neuroscience, 3*(5), 489-495.

Simons, D. J. & Chabris, C. F. (1999). Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception, 28*, 1059-1074.

Simons, D. J. & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences, 1*(7), 261-267.

Smithson, H. & Mollon, J. (2006). Do masks terminate the icon? *The Quarterly Journal of Experimental Psychology, 59*(1), 150-160.

Snowden, R. J., Willey, J., Muir, J. L. (2001). Visuospatial attention: the role of target contrast and task difficulty when assessing the effects of cues. *Perception, 30*, 983-991.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs, 74*(11), 1-29.

Sperling, G., Melchner, M. J. (1978). The attention operating characteristic: examples from visual search. *Science, 202*, 315-318.

Strasburger, H. (2005). Unfocussed spatial attention underlies the crowding effect in indirect form vision. *Journal of Vision, 5*, 1024-1037.

Thornton, I. M. & Hubbard, T. L. (2002). Representational momentum: new findings, new directions. *Visual Cognition, 9*(1/2), 1-7.

Toet, A. & Levi, D. M. (1992). The two-dimensional shape of spatial interaction zones in the parafovea. *Vision Research, 32*, 1349-1357.

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

Treisman, A. & Gormican, S. (1988). Feature analysis in early vision: evidence from search asymmetries. *Psychological Review, 95*(1), 15-48.

Tripathy, S. P. & Barrett, B. T. (2003a). Last but not least. *Perception, 32*, 1403-1408.

Tripathy, S. P. & Barrett, B. T. (2003b). Severe loss of positional information when tracking multiple dots [abstract]. *Journal of Vision, 3*(9), 340a.

Tripathy, S. P. & Barrett, B. T. (2004). Severe loss of positional information when detecting deviations in multiple trajectories. *Journal of Vision, 4*, 1020-1043.

Tripathy, S. P. & Cavangh, P. (2002). The extent of crowding in peripheral vision does not scale with target size. *Vision Research, 42*, 2357-2369.

Tripathy, S. P., Narasimhan, S. & Barrett, B. T. (2007). On the effective number of tracked trajectories in normal human vision. *Journal of Vision, 7*, 1-18.

Tsal, Y. & Bareket, T. (1999). Effects of attention on localization of stimuli in the visual field. *Psychological Bulletin & Review, 6*, 292-296.

Ungerleider, L. G. & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, Goodale, M. A. & Mansfield, R. J. W. (Ed.), *Analysis of visual behavior*. Cambridge, MA: MIT Press.

Van der Lubbe, R. H. J. & Keuss, P. J. G. (2001). Focussed attention reduces the effect of lateral interference in multi-element arrays. *Psychological Research, 65*, 107-118.

vanMarle, K. & Scholl, B. J. (2003). Attentive tracking of objects versus substances. *Psychological Science, 14*(5), 498-504.

VanRullen, R., Carlson, T. & Cavanagh, P. (2007). Dividing attention between multiple targets: simultaneous or sequential allocation? [abstract]. *Journal of Vision, 7*(9), 642a.

Ward, R., Duncan, J. & Shapiro, K. (1996). The slow time-course of visual attention. *Cognitive Psychology, 30*, 79-109.

Weichselgartner, E. & Sperling, G. (1987). Dynamics of automatic and controlled visual attention. *Science, 238*, 778-780.

Werkhoven, P., Snippe, H. P. & Toet, A. (1992). Visual processing of optic acceleration. *Vision Research, 32*(12), 2313-2329.

Wheeler, M. E. & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General, 131*(1), 48-64.

Wilken, P. & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision, 4,* 1120-1135.

Wilkinson, F., Wilson, H. R. & Ellemberg, D. (1997). Lateral interactions in peripherally viewed texture arrays. *Journal of the Optical Society of America A, 14*(9), 2057-2068.

Wolfe, J. M., Alvarez, G. A. & Horowitz, T. S. (2000). Attention is fast but volition is slow. *Nature, 406*, 691.

Xu, Y. (2002). Limitations of object-based feature encoding in visual short-term memory. *Journal of Experimental Psychology: Human Perception & Performance, 28*(2), 458-468.

Yantis, S. (1992). Multielement visual tracking: attention and perceptual organisation. *Cognitive Psychology, 24*, 295-340.

Yantis, S. (1993). Simulus-driven attentional capture and attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance, 19*(3), 676-681.

Yantis, S. & Johnson, D. N. (1990). Mechanisms of attentional priority. *Journal of Experimental Psychology: Human Perception and Performance, 16*(4), 812-825.

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

APPENDICES

APPENDIX 1: INDIVIDUAL ERROR DISTRIBUTION PLOTS

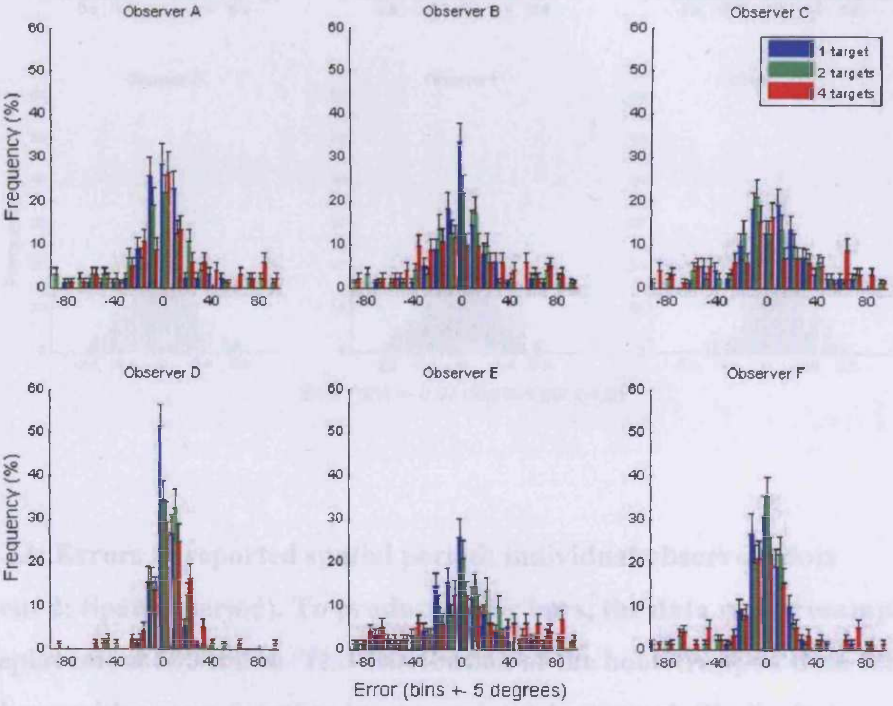

Figure 32: Errors in reported orientation: individual observer plots (Experiment 1: Orientation). There are some differences between individual observers in terms of accuracy and the effect of the number of targets for tracking on performance. Observer D, for instance, shows a high degree of accuracy and a clear effect of set size. In contrast, Observers A and C show less set size effect, but less overall precision than Observer D. For the plots shown here and also for Figures 33 and 34, error bars represent 95% confidence intervals obtained by bootstrapping.

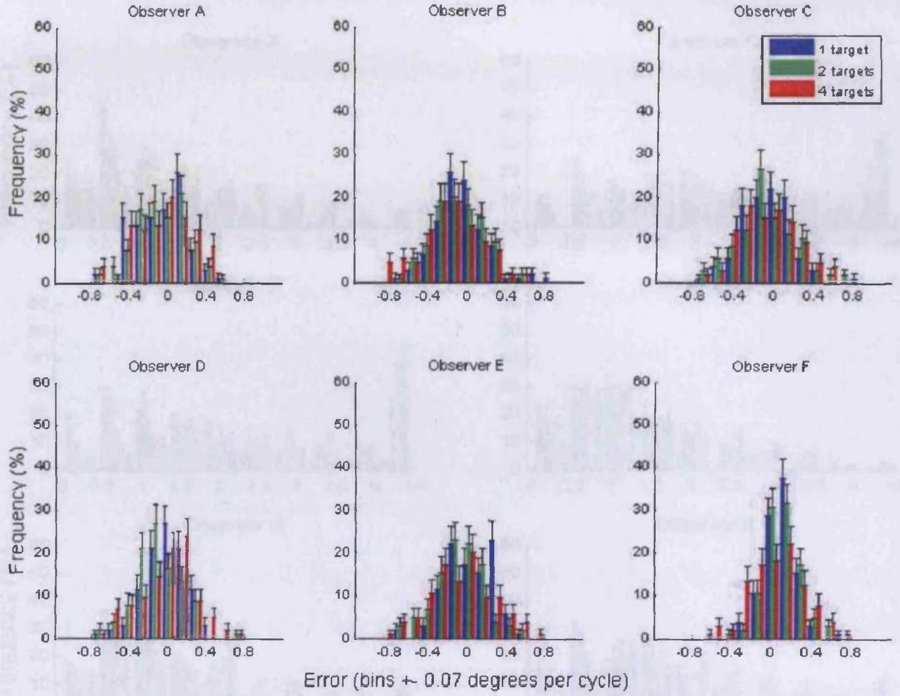

Figure 33: Errors in reported spatial period: individual observer plots (Experiment 2: Spatial period). To produce error bars, the data were resampled without replacement 500 times. The distribution of the bootstrapped data for individual error bins were found to be approximately normal. Similarly to Experiment 1, there are some differences between observers in terms of precision and set size effects. Observer F shows particularly high levels of precision and corresponding near-zero error magnitudes, especially for the tracking one and tracking two conditions. Error bars represent 95% confidence intervals obtained by bootstrapping.

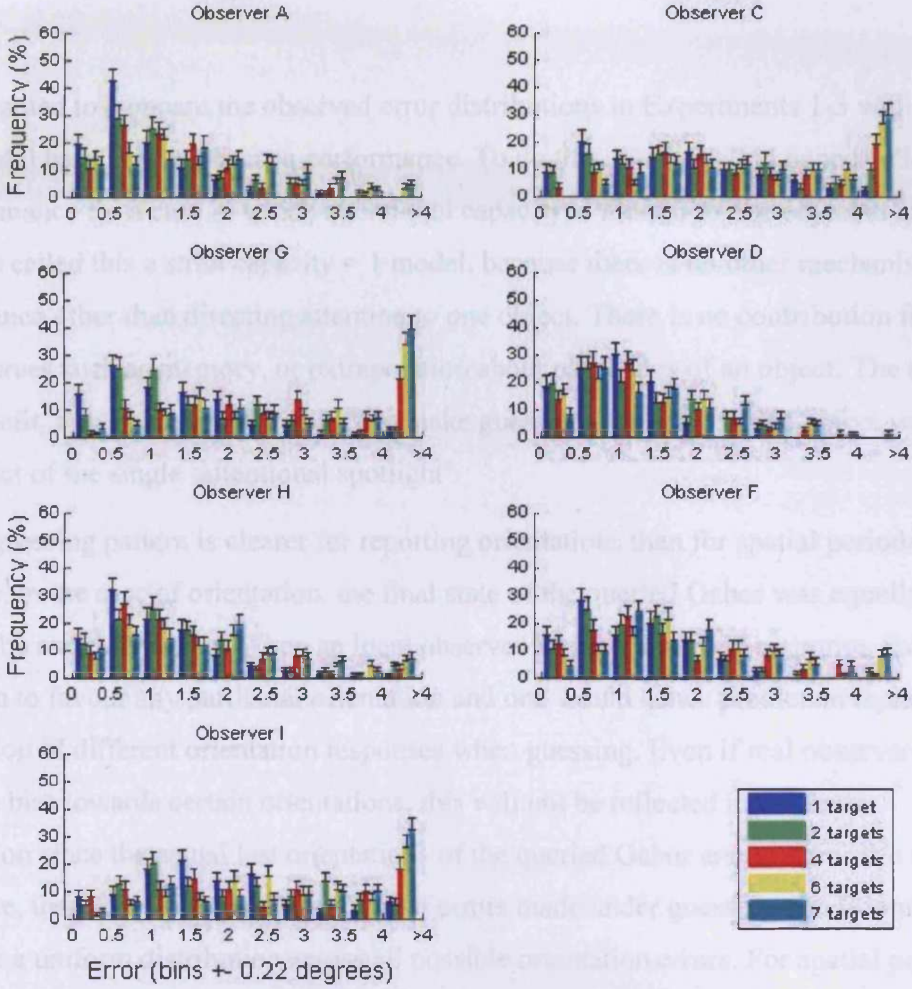

Figure 34: Errors in reported position: individual observer plots (Experiment 3: Spatial position). As in Experiments 1 and 2, some individual differences are observed. Overall, however, there is a general trend towards greater levels of spatial imprecision for greater set sizes. Error bars represent 95% confidence intervals obtained by bootstrapping.

APPENDIX 2: PREDICTIONS FROM A STRICT CAPACITY = 1 OBJECT MODEL

We wanted to compare the observed error distributions in Experiments 1-3 with a more useful baseline than chance performance. To do this, we modelled expected levels of performance for a case in which attentional capacity is limited to one object at any one time. We called this a strict capacity = 1 model, because there is no other mechanism for performance other than directing attention to one object. There is no contribution from other sources such as memory, or extrapolation about past states of an object. The model does benefit, however, from the ability to make guesses when the queried object was not the subject of the single 'attentional spotlight'.

This guessing pattern is clearer for reporting orientations than for spatial periods or positions. In the case of orientation, the final state of the queried Gabor was equally likely to be any orientation. When an ideal observer makes a guessing response, there is no reason to favour any particular orientation and one would hence predict an equal distribution of different orientation responses when guessing. Even if real observers do display a bias towards certain orientations, this will not be reflected in the error distribution since the actual last orientations of the queried Gabor are random. We can be fairly sure, therefore, that the distribution of errors made under guessing conditions would be a uniform distribution across all possible orientation errors. For spatial periods and positions, the pattern is not as clear, because the last states of the queried objects were not random and were clustered around a mean value. This makes it likely that the guessing distribution followed a tendency to cluster around the mean last state. Since we cannot be sure of the exact characteristics of these distributions, we were not able to model performance of the strict capacity = 1 model for spatial period or position monitoring, although the implications are discussed at the end of this appendix section.

Under the strict capacity = 1 model, for an orientation monitoring trial with n targets, each target object is 'spotlighted' by attention at any given moment with a probability of 1/n. Additionally we assume the strict criterion in which no memory of targets is retained once the 'spotlight' moves away. When a target is not currently 'spotlighted' by attention, participants are assumed to perform at chance levels according to the uniform guessing distribution described above. When a target is currently spotlighted by attention, performance follows that of the error distribution measured for monitoring one object. For monitoring two objects, the total modelled distribution will be the sum of the equal probabilities of the queried object being monitored (and hence following the monitoring one object distribution) and of the queried object not being monitored (and hence following the chance or guessing distribution). For monitoring four objects, the total modelled distribution will be the sum of the 25% probability that the queried object was monitored, and the 75% probability of guessing.

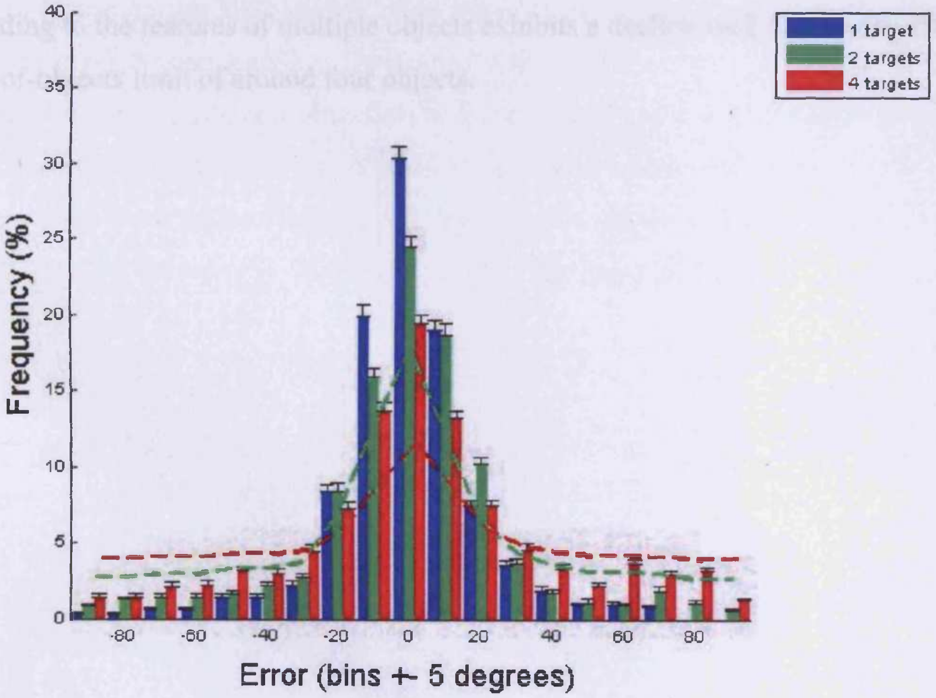

Figure 35: Predictions of a strict capacity = 1 model for monitoring two objects (green dashed line) and for monitoring four objects (red dashed line) superimposed on error distribution data from Experiment 1 (Figure 7).

Real human performance is better than that predicted by this strict capacity = 1 model and Levene's tests for equality of variance of the error distributions confirmed that model performance is significantly worse than the human data (difference between the model and human performance for monitoring two Gabors: Levene statistic = 70.001, $p < 0.01$, for monitoring four Gabors: Levene statistic = 59.904, $p < 0.01$).

For spatial period and position monitoring, it seems highly probable that the guessing strategies employed would be clustered at least to some extent around the most frequent last state of the queried object. For this reason, we might expect that the capacity = 1 model would make predictions closer to the observed levels of performance than are seen for the model comparison with orientation monitoring. Thus we would expect a greater proportion of near-zero errors even when purely guessing for spatial period and position than orientation. This would only strengthen further the conclusion that the capacity limit for attending to the features of multiple objects exhibits a decline well before any fixed-number-of-objects limit of around four objects.

APPENDIX 3: PILOT EXPERIMENTS WITH CONSTANTLY VISIBLE TARGET MARKER CUES AND WITH NO NON-TARGETS

We ran two pilot experiments for both the orientation and spatial period monitoring experiments. The first was a version where the target marker cues stayed on throughout the presentation of the Gabors to eliminate the possibility that the observed set size was due to forgetting which Gabors were the target(s). In this version, we had to increase the size of the post-cue from 0.48 degrees to 0.73 degrees to prevent the offset of the black target marker, measuring 0.48 degrees, from masking the appearance of the white post-cue appearing at the same location. The second was a version where no non-targets were ever displayed but was otherwise identical to Experiments 1 and 2. This was to ensure that the observed set size effect was not due to the distracting effects of exogenous attention capture from the non-targets. All four experiments were run with three observers, each participating in 70 trials for each of the three set sizes.

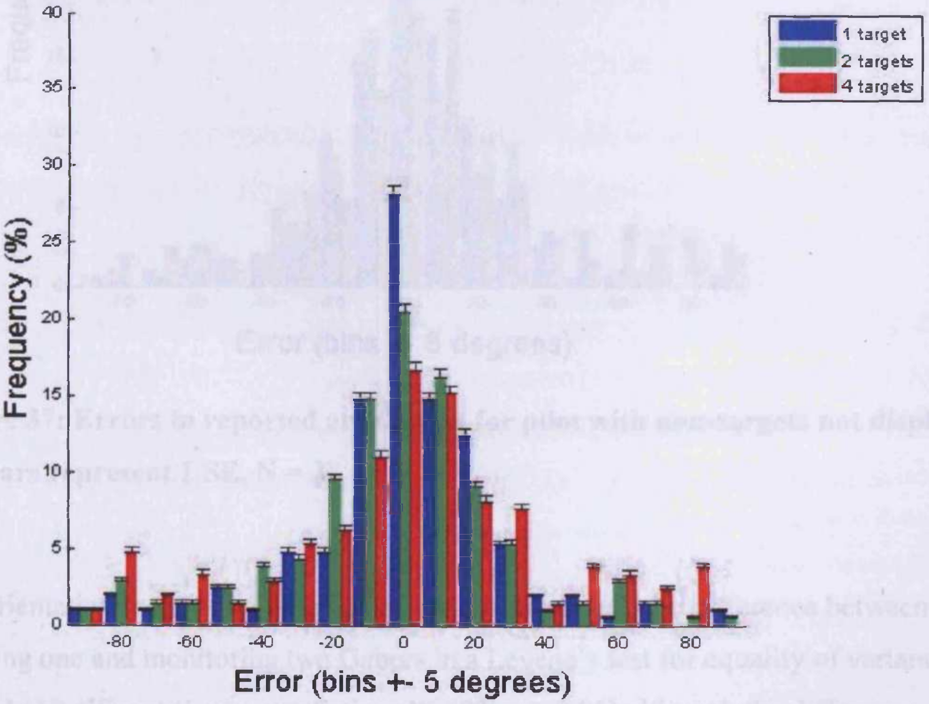

Figure 36: Errors in reported orientation for pilot with target marker cues visible throughout Gabor presentation (error bars represent 1 SE, N = 3).

For orientation monitoring with cues visible throughout the trial, there was no significant difference between monitoring one and monitoring two Gabors in a Levene's test for equality of variance (Levene statistic = 2.856, p = 0.092) although the difference between monitoring two and monitoring four was significant (Levene statistic = 7.335, p < 0.01).

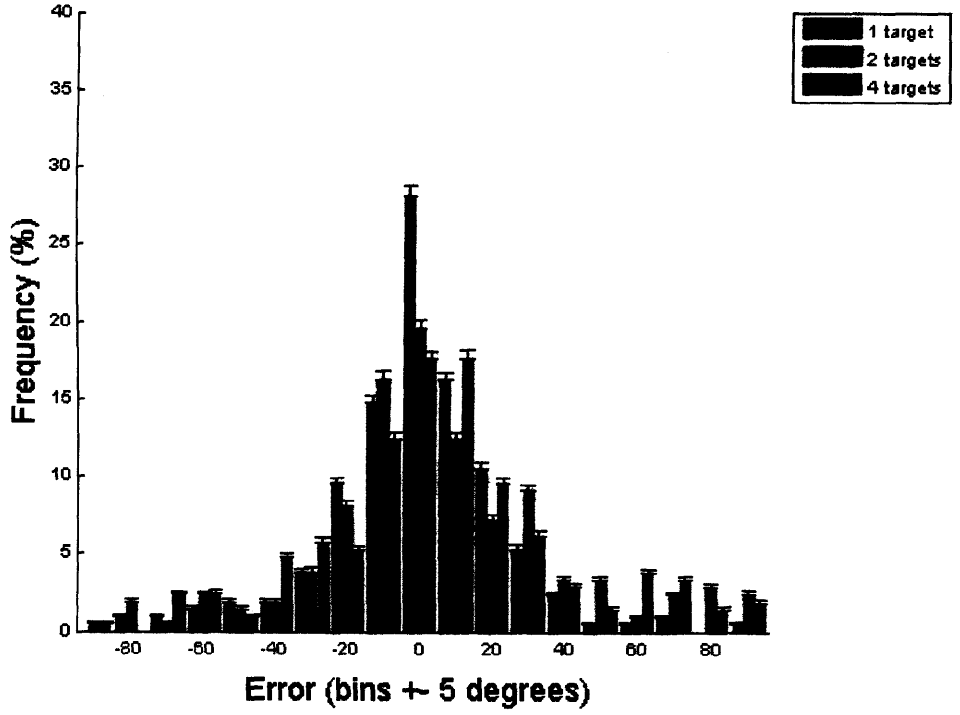

Figure 37: Errors in reported orientation for pilot with non-targets not displayed (error bars represent 1 SE, N = 3).

For orientation monitoring with no non-targets displayed, the difference between monitoring one and monitoring two Gabors in a Levene's test for equality of variance was highly significant (Levene statistic = 16.482, p < 0.01) although the difference between monitoring two and monitoring four was not significant (Levene statistic = 0.0, p = 0.988).

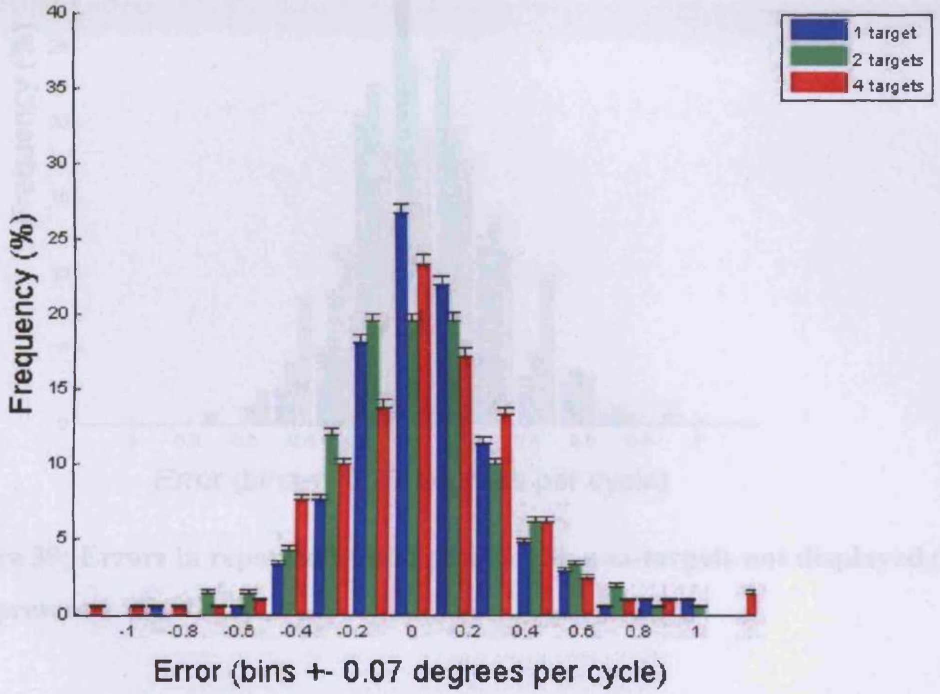

Figure 38: Errors in reported spatial period for pilot with target marker cues visible throughout Gabor presentation (error bars represent 1 SE, N = 3).

For spatial period monitoring with cues visible throughout Gabor presentation, the difference between monitoring one and monitoring two Gabors in a Levene's test for equality of variance was significant (Levene statistic = 5.560, p < 0.05) although the difference between monitoring two and monitoring four was not (Levene statistic = 0.264, p = 0.608).

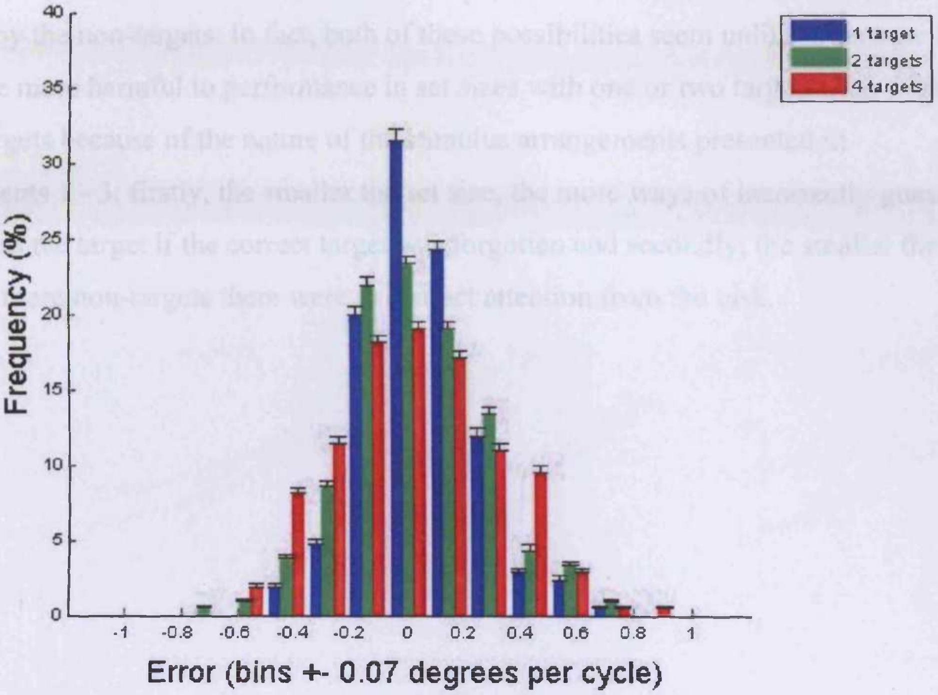

Figure 39: Errors in reported spatial period with non-targets not displayed (error bars represent 1 SE, N = 3).

For spatial period monitoring with no non-targets displayed, the difference between monitoring one and monitoring two Gabors in a Levene's test for equality of variance was highly significant (Levene statistic = 8.647, $p < 0.01$) and the difference between monitoring two and monitoring four was also significant (Levene statistic = 5.543, $p < 0.05$).

In these four pilot experiments, although some of the differences between set sizes were not significant, each experiment had a significant difference in variance of errors between at least either monitoring one and two objects, or between monitoring two and four objects. Each observer only participated in two-thirds of the total number of trials as in the main Experiments 1-3, and there were only three observers in each experiment: the data in these four experiments are highly suggestive that with more power, we might find similar magnitudes of set size effects as observed in Experiments 1 and 2. At the very least, it does not appear that the set size effects observed in Experiments 1 and 2 were due to either forgetting about which Gabors were targets, or due to exogenous attentional capture by the non-targets. In fact, both of these possibilities seem unlikely because both would be more harmful to performance in set sizes with one or two targets than with very many targets because of the nature of the stimulus arrangements presented in Experiments 1 - 3: firstly, the smaller the set size, the more ways of incorrectly guessing which was the target if the correct target was forgotten and secondly, the smaller the set size, the more non-targets there were to distract attention from the task.

APPENDIX 4: INDIVIDUAL LAG PLOTS

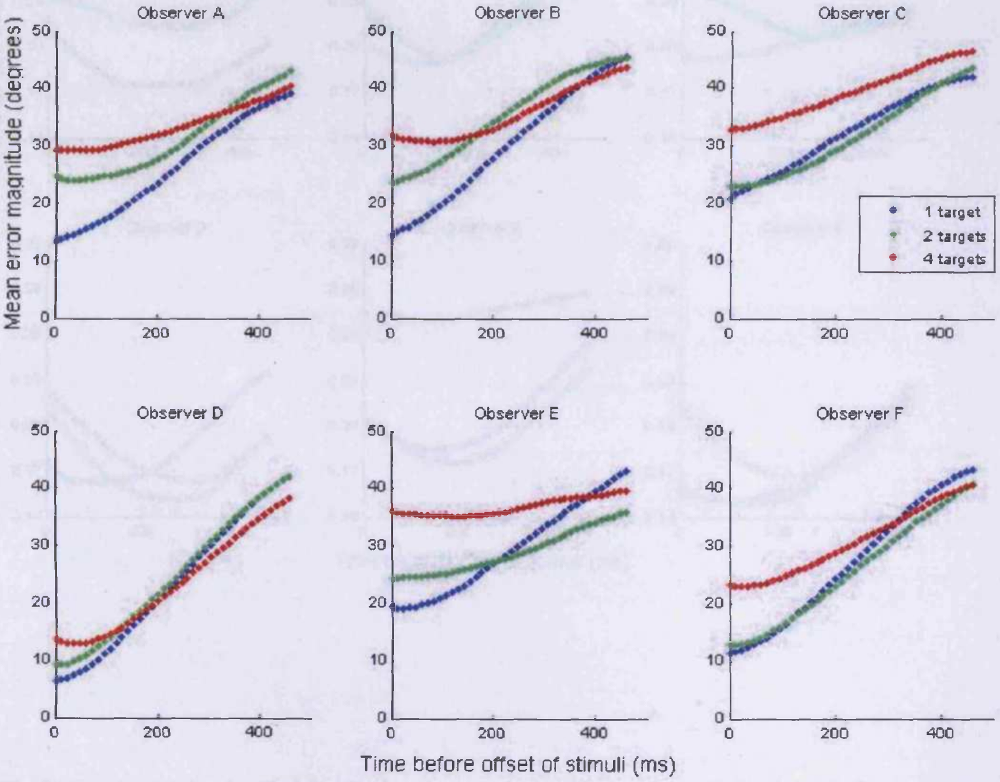

Figure 40: Lag analysis: individual observer plots (Experiment 1: Orientation). Different observers display different overall mean error magnitudes and different effects of set size on precision. For instance, Observer D shows a relatively high level of precision indicated by the low mean error magnitudes for the leftmost points. In addition this observer shows relatively little effect of the number of targets for tracking on precision. These plots also illustrate perceptual lag magnitudes, and variability in perceptual lags with set size. For instance, Observer D shows a clear

increase in perceptual lags with increases in set size. Conversely, Observer C does not exhibit large perceptual lags, and shows little increase in lags with set size.

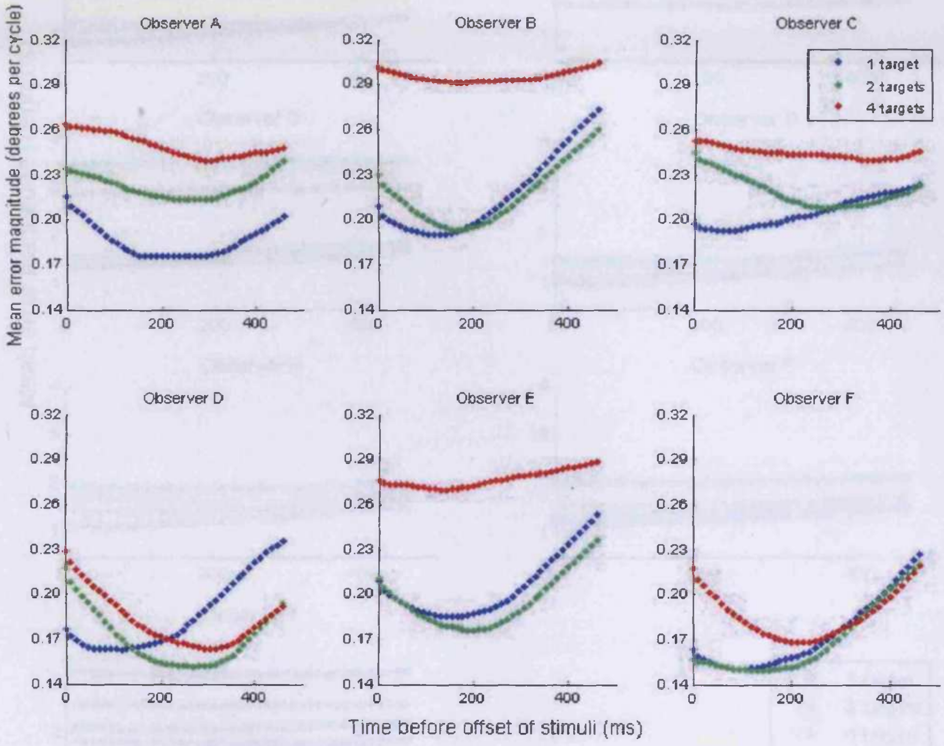

Figure 41: Lag analysis: individual observer plots (Experiment 2: Spatial period). Large individual differences are observed in perceptual lag patterns: Observer A shows an increase in mean error magnitude with set size, but little effect of set size on perceptual lags. At the other extreme, it appears that the apparent increase in error magnitude with set size shown by Observer D was in fact mostly contributed to by an increase in perceptual lag with set size.

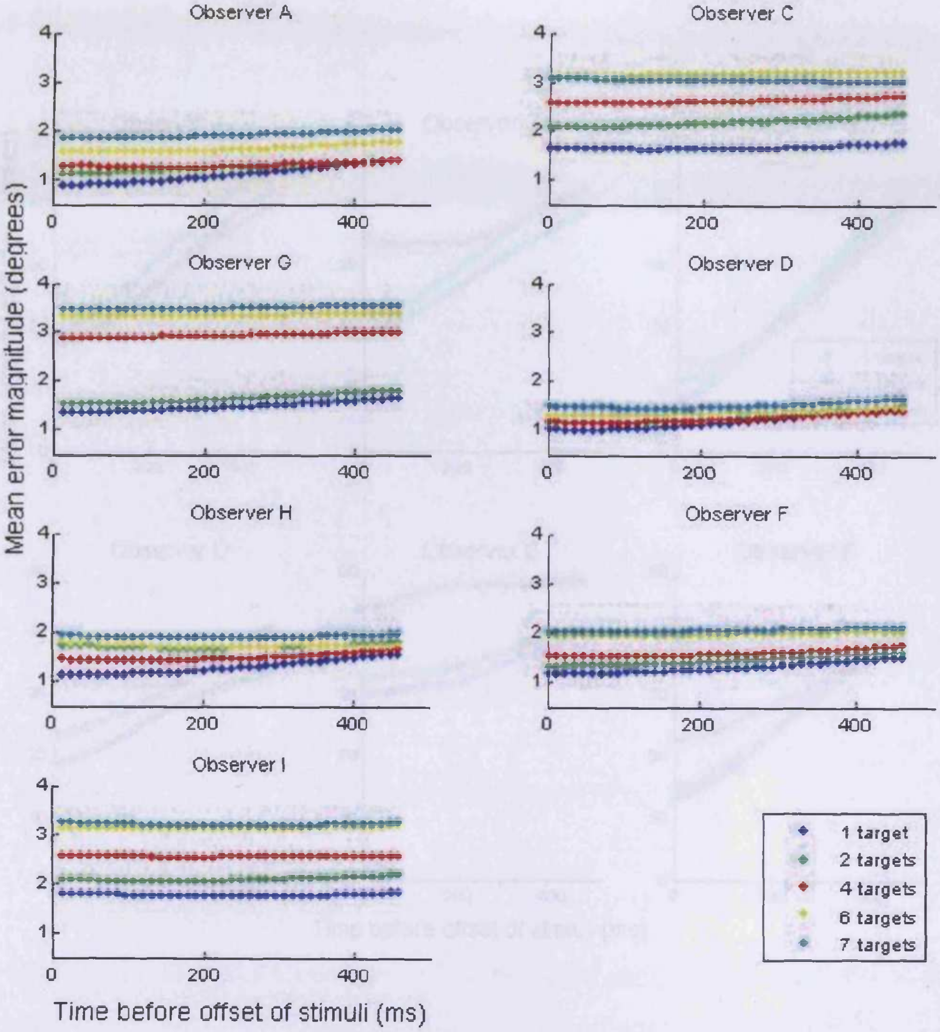

Figure 42: Lag analysis: individual observer plots (Experiment 3: Position). Some individual differences are observed for position tracking. There are differences in overall error magnitudes and in set size effects on precision. Some variability in lag patterns is also observed.

APPENDIX 5: INDIVIDUAL LAG PLOTS FOR EXPERIMENTS 6 - 9

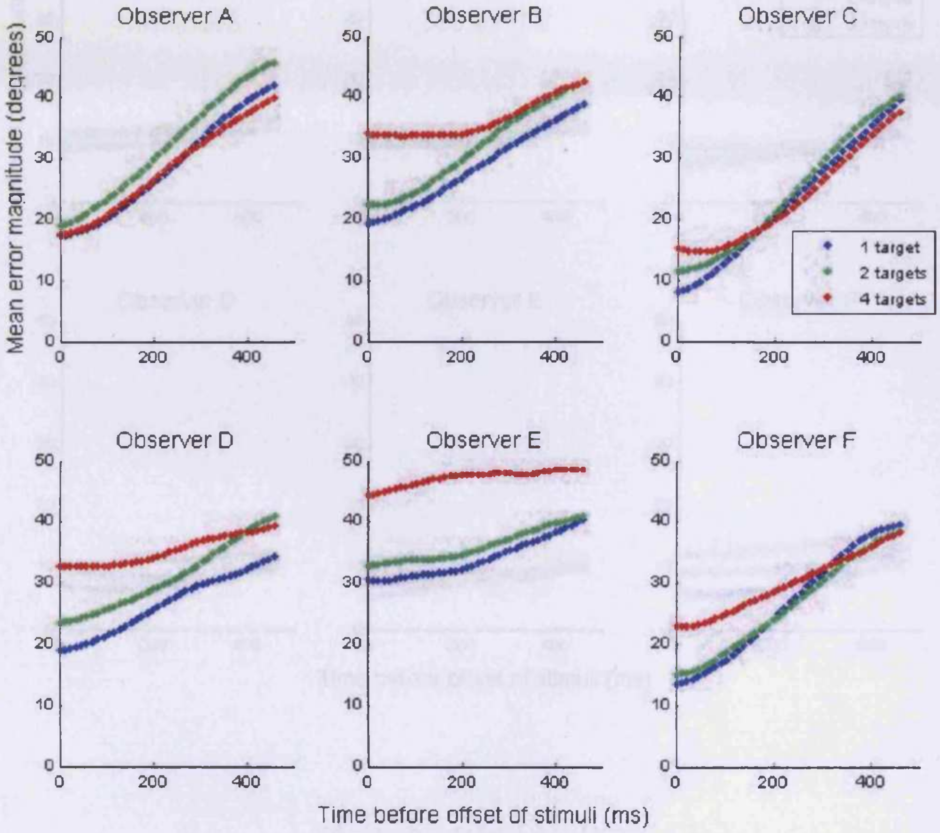

Figure 43: Lag analysis for Experiment 6: individual observer plots. This experiment included sharp reversals when orientations reach 0 or 360 degrees. There are differences in overall error magnitudes, lags and in set size effects. There is no evidence, however, of the very large lags and large set size effects observed in Experiment 2.

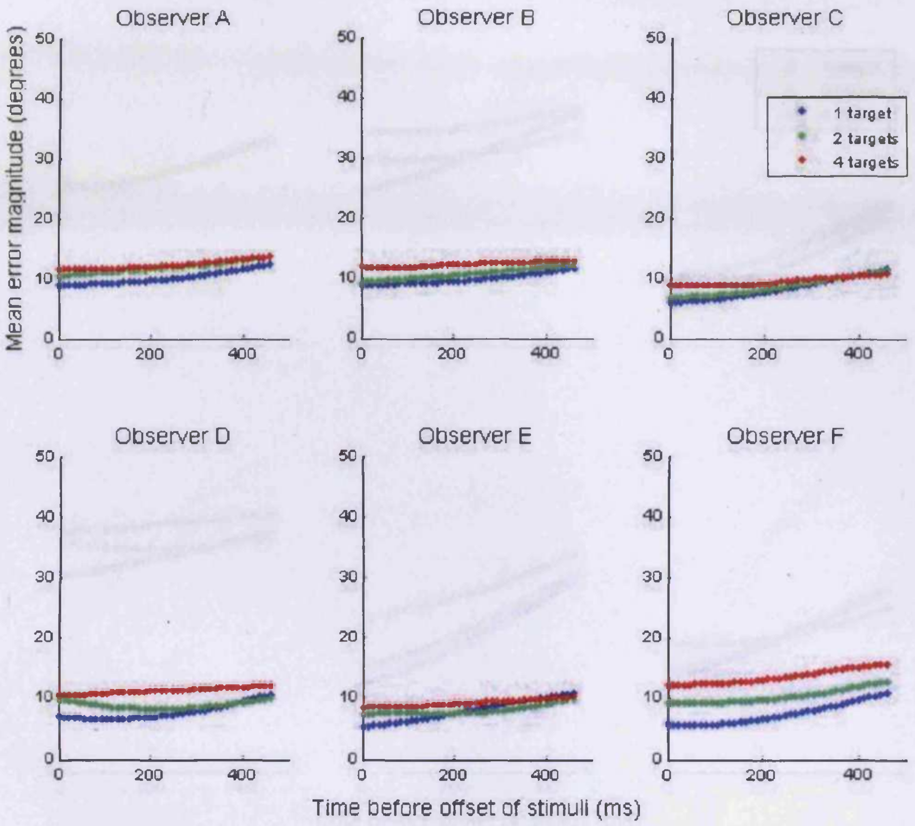

Figure 44: Lag analysis for Experiment 7: individual observer plots, for which the orientation range was restricted to between 0 and 45 degrees. Similarly to Experiment 6, there are differences in overall error magnitudes, lags and in set size effects. There is no evidence, however, of the very large lags and large set size effects observed in Experiment 2.

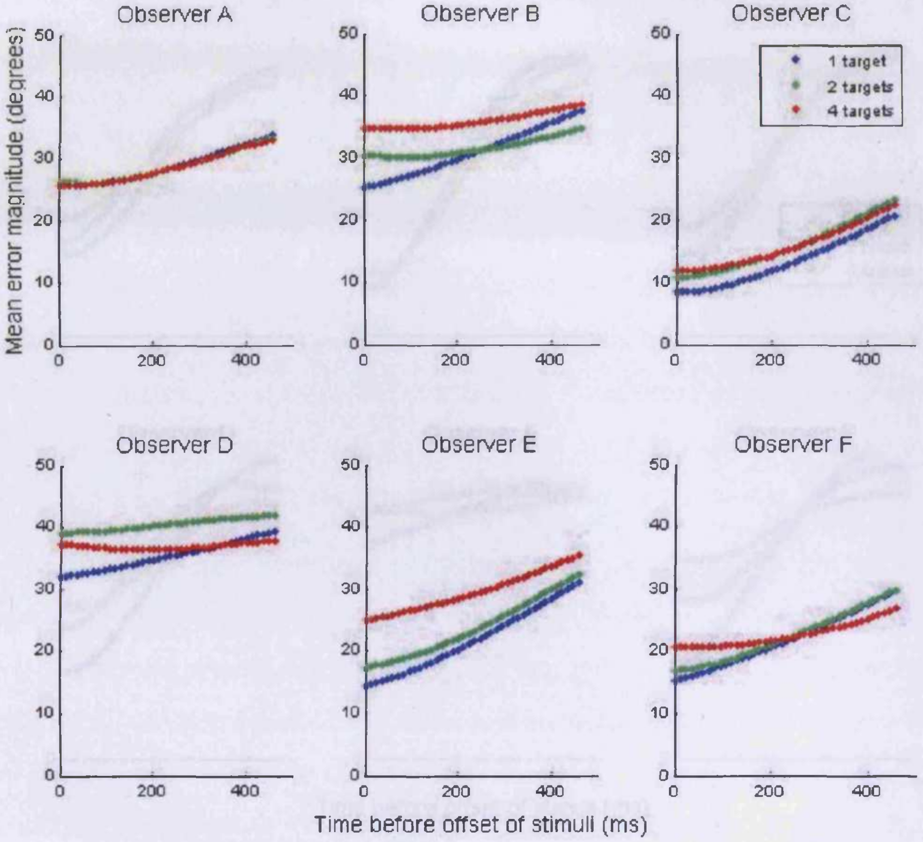

Figure 45: Lag analysis for Experiment 8: individual observer plots. This experiment used slow rates of orientation change. Again, there are differences between observers in overall error magnitudes, lags and in set size effects. There is still no evidence of the very large lags and large set size effects observed in Experiment 2.

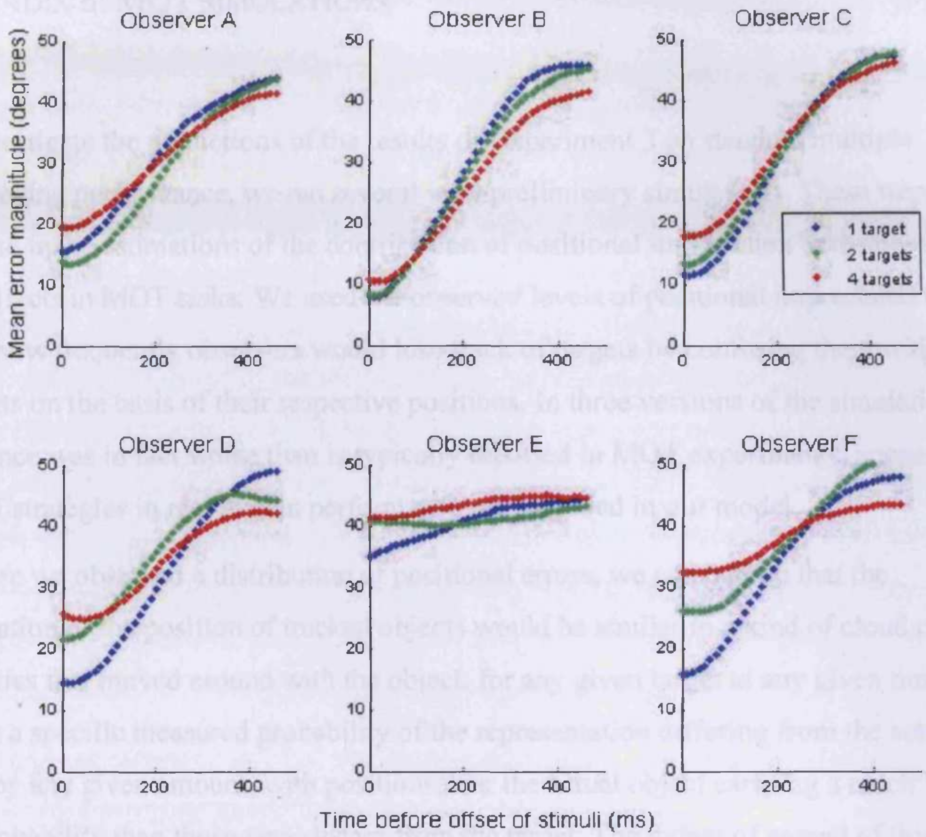

Figure 46: Lag analysis for Experiment 9: individual observer plots. This experiment used fast rates of orientation change. Differences between observers can again be seen in overall error magnitudes, lags and in set size effects. Despite the more pronounced pattern of lags, i.e. the difference in error magnitudes at the lag time and at other times, there is still no evidence of the very large lags and large set size effects observed in Experiment 2.

APPENDIX 6: MOT SIMULATIONS

To investigate the predictions of the results of Experiment 3 on standard multiple object tracking performance, we ran several very preliminary simulations. These were intended to make estimations of the contribution of positional imprecision to the observed set size effects in MOT tasks. We used the observed levels of positional imprecision to estimate how frequently observers would lose track of targets by confusing them with non-targets on the basis of their respective positions. In three versions of the simulation, performance was in fact worse than is typically reported in MOT experiments, suggesting the use of strategies in real human performance not included in our model.

Because we observed a distribution of positional errors, we considered that the representation of the position of tracked objects would be similar to a kind of cloud of probabilities that moved around with the object: for any given target at any given time, there was a specific measured probability of the representation differing from the actual position by any given amount, with positions near the actual object carrying a much higher probability than those very distant from the target. The extent of spread of this probabilistic cloud can be modelled using the positional errors recorded, and will be much more diffuse for tracking larger numbers of objects than for tracking only one or two, for instance.

We used these distributions to estimate probabilities of targets becoming confused with non-targets and thereby contributing to the well-documented decline in performance in MOT tasks as the number of targets is increased. We ran simulations in real time calculating the positions of targets and non-targets on a frame-by-frame basis, and calculating on each frame the likelihood of losing track of targets.

Targets and non-targets moved around in a quasi-random manner, bouncing off the square perimeter boundary and also changing frequently in their horizontal and vertical speeds and directions of motion (with motion parameters similar to those used in the original Pylyshyn and Storm (1988) study). Statistically, targets and non-targets will frequently pass close to one another, and the closer they become, the more likely it will be that the target will be falsely represented at or very near the location of the non-target, thus making target/non-target confusion possible. The calculation of target/non-target confusion probability was conducted in the model for every new set of positions of the stimuli such that longer periods of time spent close to one another resulted in a greater chance of a target and a non-target becoming confused with one another.

In the first model, "probabilistic tracking", whenever this type of mislocalisation occurred, the model simply selected the target or the non-target that were involved with 50% probability each, and then subsequently tracked the selected object (which may be the target or the non-target). In the second model, "spreading activation", the model tracked both of the objects involved. In this case, one object was added to the tracking load after every potential mislocalisation event. Because we know that the more objects are tracked, the greater the spread of errors, each of these mislocations caused all objects to be tracked with less precision. In the third model, "losing track", both objects involved were simply dropped from the tracking. In this case, each mislocalisation event caused an increase in precision with which the remaining targets were tracked.

Here, as in the original Pylyshyn and Storm study (1988), performance was measured in each trial by a probe on one of the targets at a semi-random interval several seconds after the start of the trial. The probe was a square covering one of the targets and observers simply had to press a button to indicate that they had detected a probe on one of the targets (squares were also frequently presented over non-targets to prevent simply to responding to any probe). In the probabilistic tracking model, a 'response' was judged to be correct simply if the probed object was being tracked at the time of the probe. In the other two models, a correct response was recorded if the probed object was being tracked. However, these two models also entailed the possibility of a correct guess. In these models, the number of objects tracked can be different from the number of targets, and the models are able to make informed guesses accordingly. The spreading activation model, for instance, when tracking three objects in a two-target trial, will guess with 67% probability that any one of the tracked objects is in fact a target.

In the Pylyshyn and Storm study, performance for tracking one target was at 96% and for tracking five targets at 86%. All three models described here performed worse than this, as shown in Table 3 below.

(% correct)	One target	Two targets	Four targets	Five targets
Human data	96	-	-	86
Probabilistic tracking model	26	25	63	-
Spreading activation model	35	54	-	-
Losing track model	6	7	20	-

Table 3: Performance of the three MOT simulation models compared to human data (Pylyshyn & Storm, 1988).

One curious outcome of the simulations is the sometimes improved performance of the models with higher numbers of targets, particularly when tracking four targets. This is contributed to by the decreased likelihood of near-collisions between targets and non-targets when there is a higher proportion of targets to non-targets. Future modelling might attempt to account for this factor.

Because we only have positional error data for tracking one, two or four targets, performance when tracking three targets was estimated by interpolating performance between tracking two and tracking four targets. Also note that no prediction can be made for the spreading activation model with four targets, since we only had positional error data available for up to four targets, and hence were unable to simulate performance when five objects were tracked.

Simulated performance was much worse than real human performance data. One possibility is of course that we cannot know whether tracking performance draws on a better level of precision than the level tapped by explicit report. With this caveat, the data point towards the use in MOT tasks of strategies such as speed or velocity monitoring in addition to position monitoring to recover lost targets and perhaps to disambiguate targets from non-targets. Indeed very recently, Horowitz et al. (2007) have shown a progressive increase in thresholds for direction of motion discrimination in MOT tasks as the number of targets is increased. This indicates that direction of motion can indeed be encoded during the task. Whether in fact it is used when it is not explicitly explained to observers as being a requirement of the task, is still however, an open question.

APPENDIX 7: DUAL TASK PILOT EXPERIMENTS

Two pilot experiments using a similar dual task method to that used in Experiment 10 yielded similar results: there was no clear pattern in the effect of attention on either sampling or local noise estimates. The method for Experiment 12 was similar to that for Experiment 10 except that there were three Gabors instead of six, and these were presented to the left, right and above the centre of the screen and at the same distance as in Experiment 10. A second difference was that the 'jump' event in the secondary task stimulus only occurred with 50% probability, so that observers were frequently still looking for the jump event when the primary task stimuli (Gabors) appeared. The method for Experiment 13 was similar to the low attentional load (cued) condition for Experiment 11 except for two differences: only the three task-relevant Gabors appeared on each trial, and in addition the secondary task stimulus was presented before each trial. The results of these two pilot experiments are presented below.

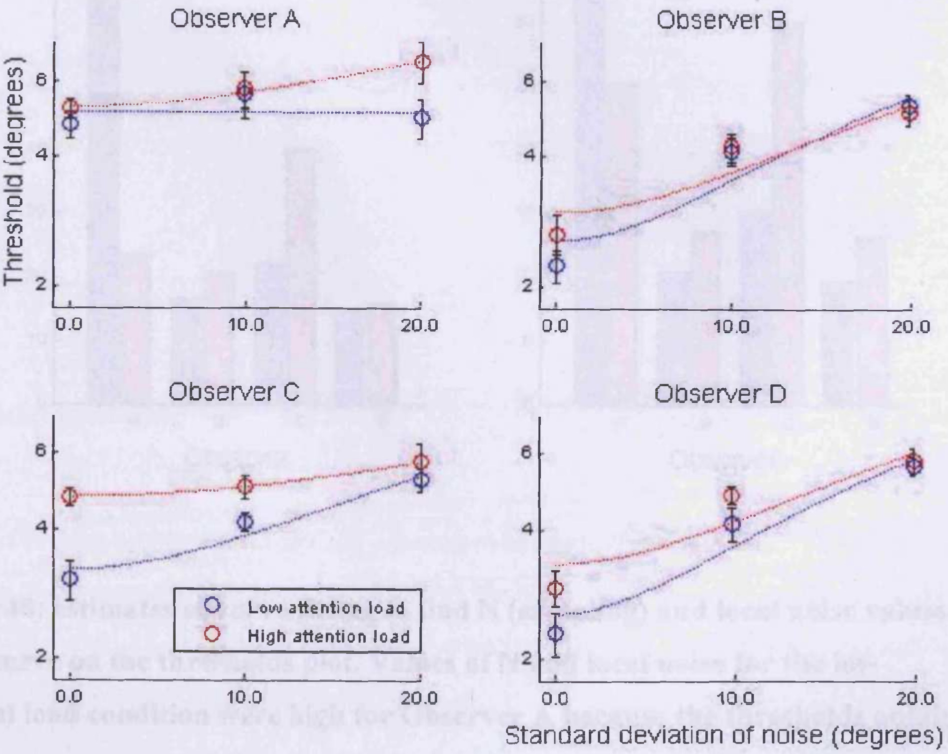

Figure 47: mean thresholds for four observers in each of the six total experimental conditions (3xexternal noise level and 2xattention) for Experiment 12. Error bars in black represent standard errors of the nine individual threshold measurements taken for each of the six conditions. Dotted lines represent curves fitted using the equivalent noise model.

The main effect of attention on thresholds was approaching significance (F(1, 3) = 9.190, p = 0.056). There was a significant effect of external noise level (F(2, 6) = 14.628, p < 0.01) in a 2 (attention level) x 3 (external noise level) repeated-measures ANOVA.

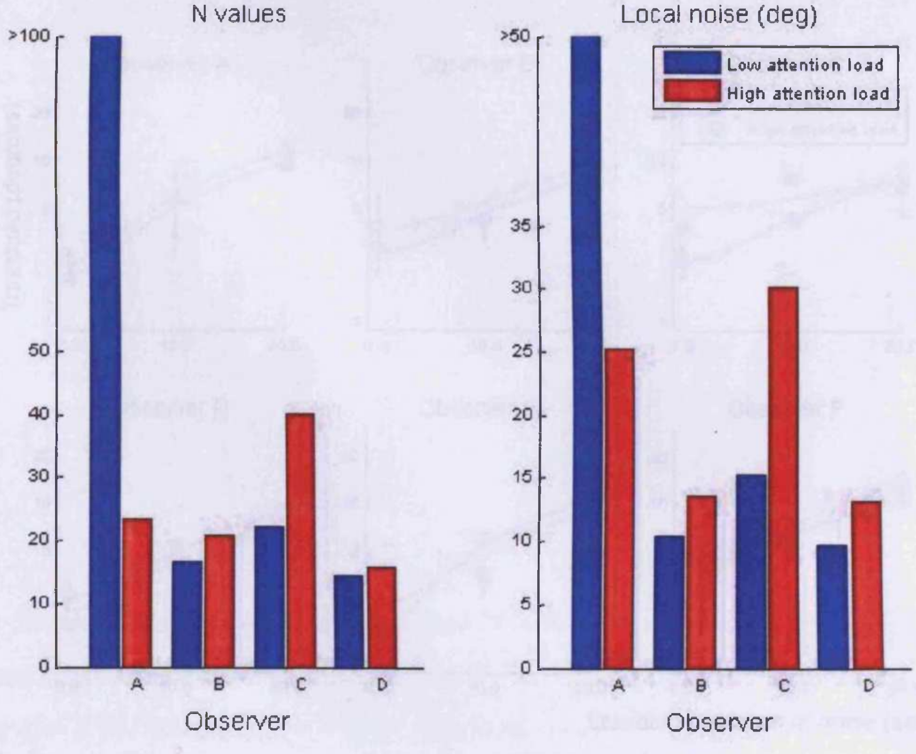

Figure 48: estimates of curve fitting to find N (sampling) and local noise values for each curve on the thresholds plot. Values of N and local noise for the low attentional load condition were high for Observer A because the thresholds obtained

violated the assumption of a monotonic increase with increasing levels of external noise. Actual fitted values were N = 6040 and local noise = 395. For this reason the pattern of attentional effects on N and local noise are not clear.

There was no significant effect of attention on N (number of samples) in a paired t-test (t(3) = 0.995, p = 0.393). Neither was there a significant effect of attention on local noise (t(3) = 0.923, p = 0.424). Neither was the effect on N (t(2) = -1.534, p = 0.265) or local noise significant (t(2) = -1.868, p = 0.203) after excluding Observer A. Indeed these results may have changed if more data were collected for more observers but as they stand the results are not conclusive.

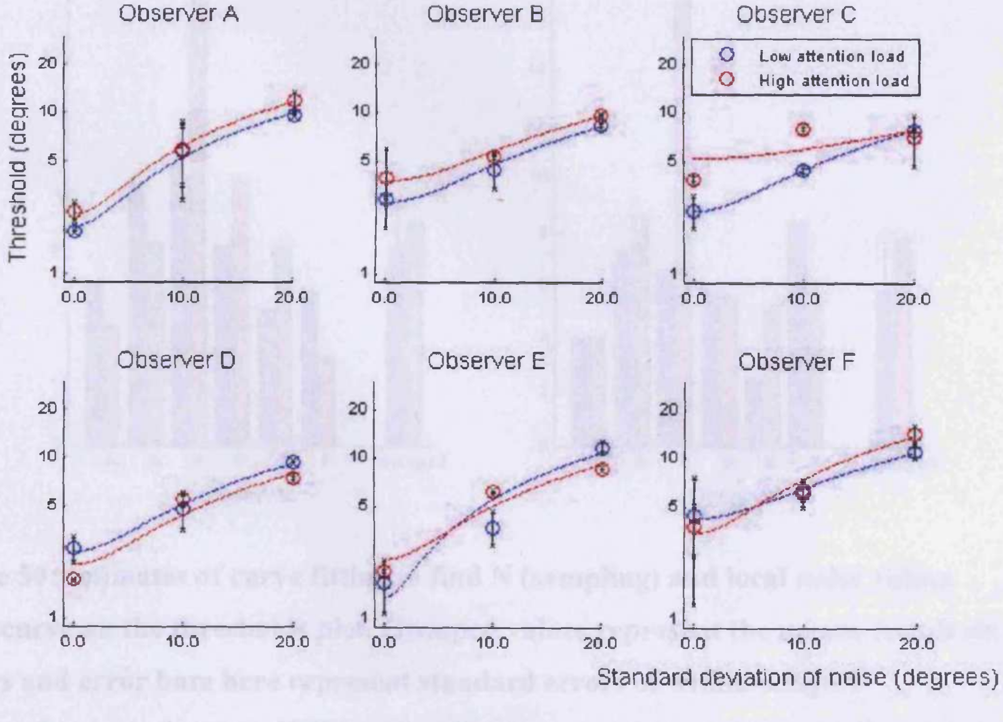

Figure 49: mean thresholds for six observers in each of the six total experimental conditions (3xexternal noise level and 2xattention) for Experiment 13. Error bars in black represent standard errors obtained by bootstrapping. Dotted lines represent curves fitted using the equivalent noise model.

The main effect of attention on thresholds was not significant overall (F(1, 5) = 2.065, p = 0.210) though there was a significant effect of external noise level (F(2, 10) = 1.572, p 0.255) in a 2 (attention level) x 3 (external noise level) repeated-measures ANOVA. The effect of attention did become significant, however, when Observer D was excluded as Observer D showed higher thresholds for the low than for high attentional load condition (F(1, 4) = 10.382, p < 0.05).

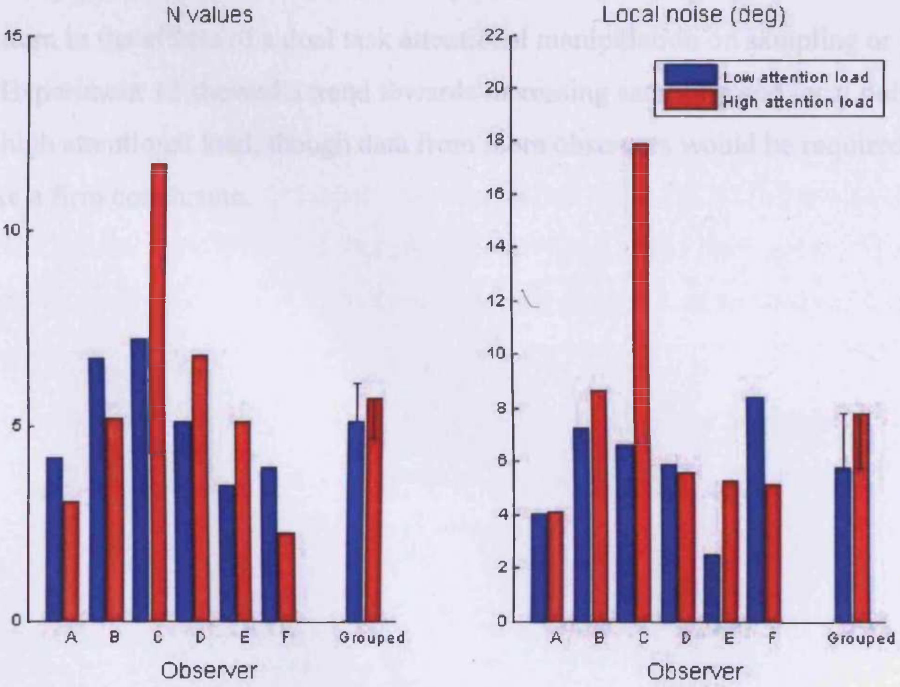

Figure 50: estimates of curve fitting to find N (sampling) and local noise values for each curve on the thresholds plot. Grouped values represent the means for all six observers and error bars here represent standard errors of within-subjects differences between the two conditions i.e. the difference between the conditions is calculated first for each observer, then the standard error of these differences is calculated. Error bars thus represent differences between conditions rather than variability within conditions.

Even excluding Observer D who displayed higher thresholds for the low than the high attentional load conditions, we can observe no clear pattern in the effect of attention on either sampling or local noise values.

There was no significant effect of attention on N (number of samples) in a paired t-test ($t(5) = -0.586$, $p = 0.583$). Neither was there a significant effect of attention on local noise ($t(5) = -0.976$, $p = 0.374$). After excluding Observer D, there was still no significant effect of attention on N (number of samples) in a paired t-test ($t(4) = -0.304$, $p = 0.776$). Neither was there a significant effect of attention on local noise ($t(4) = -1.006$, $p = 0.371$).

Although only preliminary studies, neither Experiment 12 nor Experiment 13 showed any clear pattern in the effects of a dual task attentional manipulation on sampling or on local noise. Experiment 12 showed a trend towards increasing sampling and local noise values with high attentional load, though data from more observers would be required in order to make a firm conclusion.

APPENDIX 8: ANALYSES CONDUCTED ON DATA FROM EXPERIMENTS 10 AND 11 WITH CONSTRAINTS PLACED ON THRESHOLDS

It is a plausible constraint that the measured thresholds in Experiments 10 and 11 should rise monotonically with the amount of external noise and we wondered whether imposing this constraint might reveal a clearer pattern in the N and local noise values arising from curve fitting. Only very slight differences were observed however in either experiment.

In both experiments, before every curve fit was conducted, a constraint was imposed that each set of three threshold values should obey the rule that each successive threshold should be either equal to, or greater than the last as external noise values increased. On occasions when this rule was broken, a straight line was fitted to the three points. The curve fit was then applied to the three points at which the straight line crossed the three external noise values.

Violation of the monotonic rule occurred for one observer in each experiment, in Experiment 10 for Observer C and in Experiment 11 for Observer F. The curve fits subsequently obtained are shown in Figures 51 and 52.

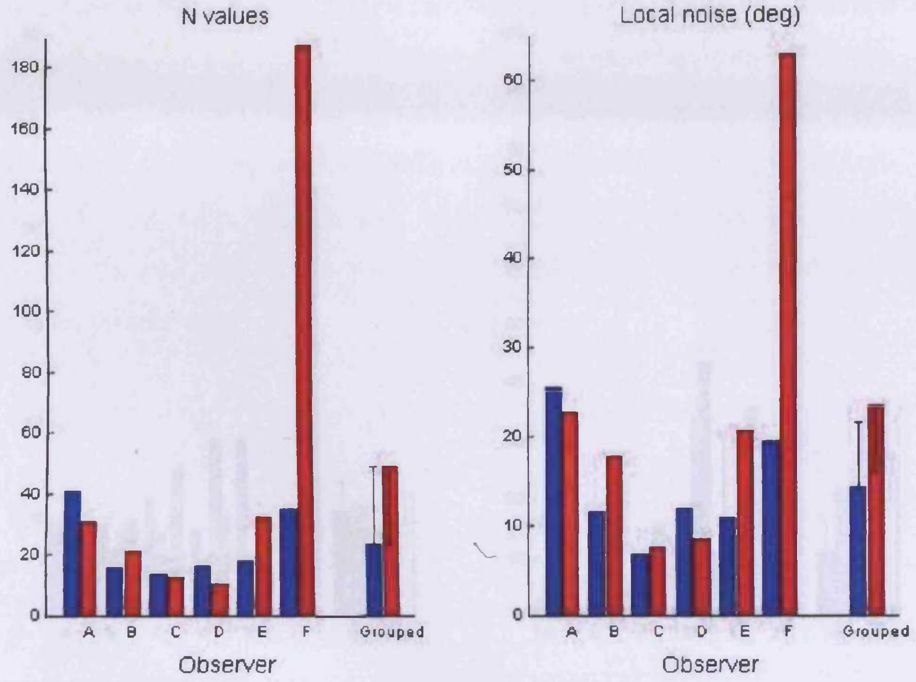

Figure 51: estimates of curve fitting to find N and local noise values after monotonic constraint imposition for Experiment 10. Grouped values represent the means for all six observers and error bars here represent standard errors of within-subjects differences between the two conditions.

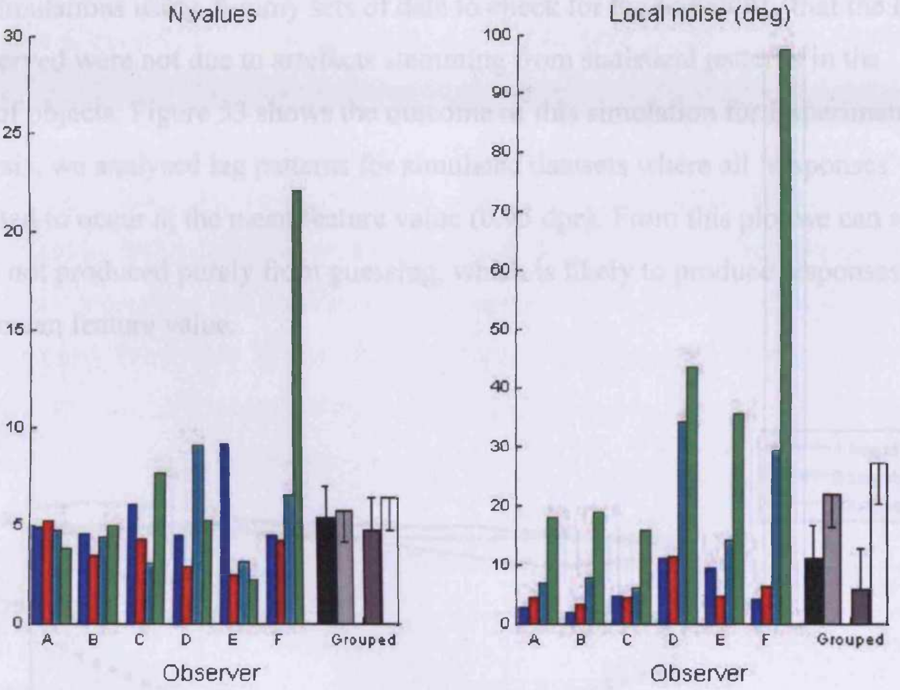

Figure 52: estimates from the results of curve fitting to find N and local noise values after monotonic constraint imposition for Experiment 11. Grouped values represent the means for all six observers for all low attention trials and all high attention trials (irrespective of crowding), all trials without crowding and all trials with crowding (irrespective of attentional load), respectively. Error bars here represent standard errors of within-subjects differences between the two attentional and crowding conditions.

APPENDIX 9: SIMULATED LAG PATTERNS

We ran simulations using dummy sets of data to check for the possibility that the lag patterns observed were not due to artefacts stemming from statistical patterns in the trajectories of objects. Figure 53 shows the outcome of this simulation for Experiment 2. In this analysis, we analysed lag patterns for simulated datasets where all ‘responses’ were simulated to occur at the mean feature value (0.95 dpc). From this plot we can see that lags are not produced purely from guessing, which is likely to produce responses near to this mean feature value.

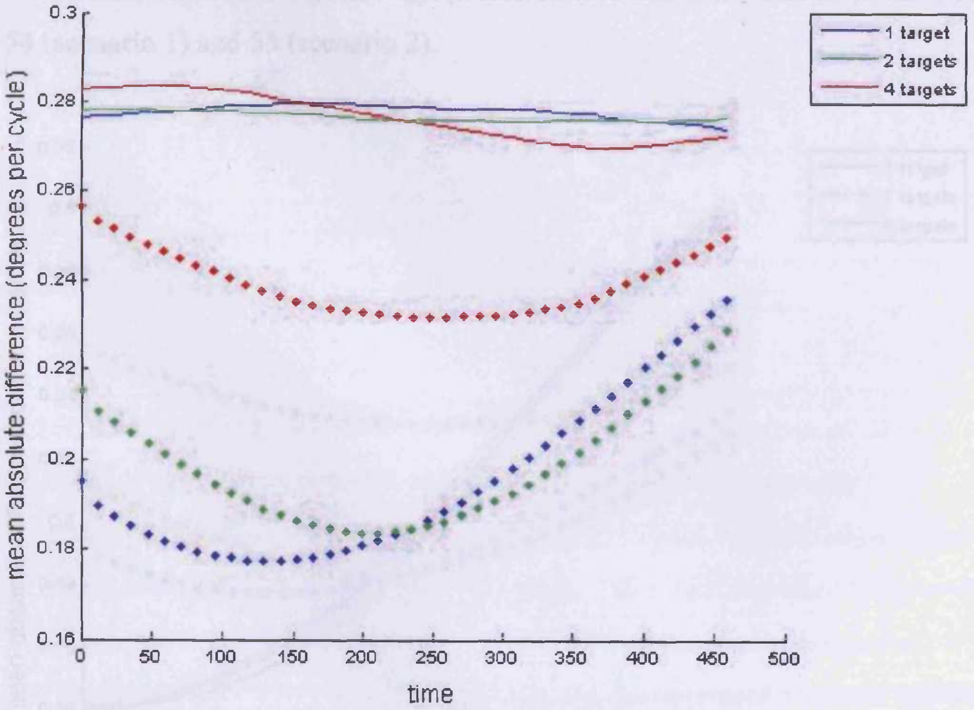

Figure 53: Simulated lags are shown in solid lines: here all ‘responses’ are simulated to occur at the mean spatial period (0.95 dpc). No lag pattern is observed in this simulated data, compared to the behavioural data previously shown in Figure 11 (indicated with dotted lines).

On those trials where the last feature value of the queried object was relatively near to this mean feature value (scenario 1), guessing will produce smaller error magnitudes and also shorter lags. Conversely, on those trials where the last feature value of the queried object is relatively far from the mean feature value (scenario 2), guessing will produce greater error magnitudes and longer lags. The latter case makes intuitive sense when we consider that the queried object is likely to have possessed a more moderate feature value at a previous time in the trial, displaying the statistical characteristics of regression to the mean. We would expect that this absence of lag observed in the guessing simulation is a product of short lags from trials in scenario 1, and long lags in scenario 2, that purely result from the statistics of the objects' trajectories. Indeed this is the case as is shown in Figures 54 (scenario 1) and 55 (scenario 2).

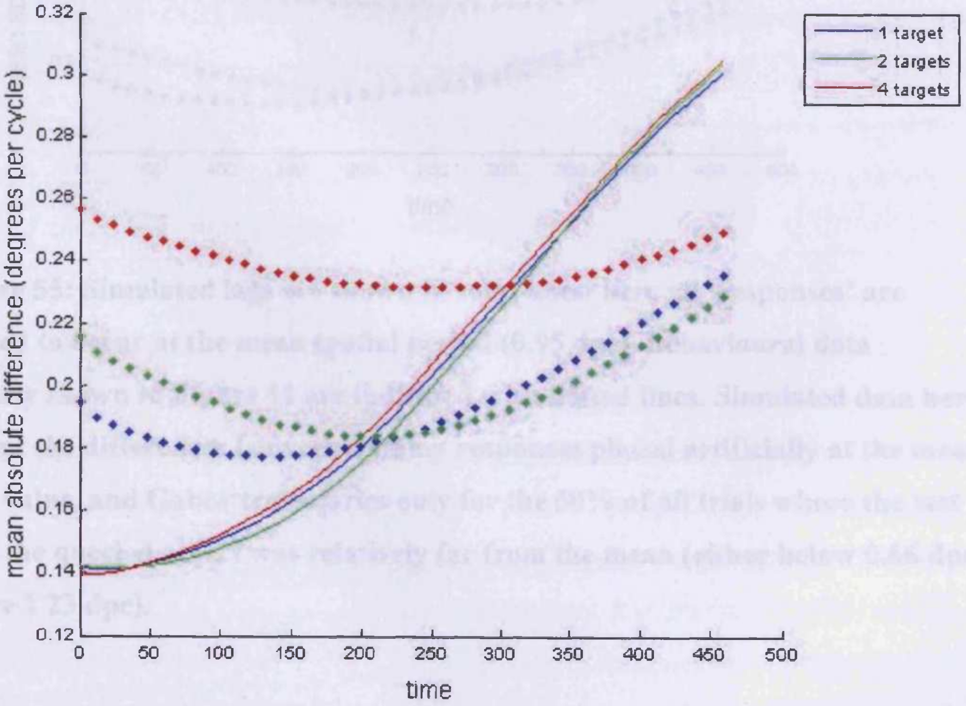

Figure 54: Simulated lags are shown in solid lines: here all 'responses' are simulated to occur at the mean spatial period (0.95 dpc). Behavioural data previously shown in Figure 11 are indicated with dotted lines. Simulated data here

represent the differences between dummy responses placed artificially at the mean feature value, and Gabor trajectories only for the 50% of all trials where the last state of the queried object was relatively moderate (between 0.66 and 1.23 dpc).

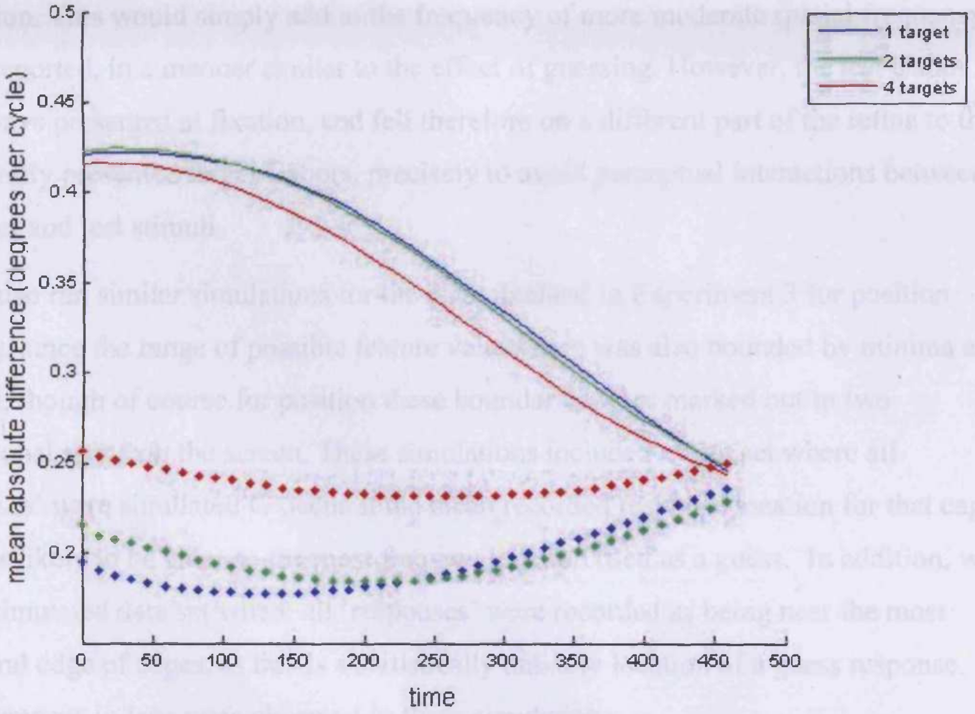

Figure 55: Simulated lags are shown in solid lines: here all 'responses' are simulated to occur at the mean spatial period (0.95 dpc). Behavioural data previously shown in Figure 11 are indicated with dotted lines. Simulated data here represent the differences between dummy responses placed artificially at the mean feature value, and Gabor trajectories only for the 50% of all trials where the last state of the queried object was relatively far from the mean (either below 0.66 dpc or above 1.23 dpc).

As is evident from Figures 54 and 55, scenarios 1 and 2 sum to produce no overall effect of a guessing strategy on lag patterns produced. In addition, there is no reason to assume that guessing should occur relatively more or less frequently in either of scenarios 1 or 2.

One might also expect that observers would make feature reports biased towards the mean spatial frequency value as a consequence of adaptation to the mean spatial frequency of Gabors (e.g. Blakemore, Nachmias & Sutton, 1970), since the test patch would appear to be more different from the mean (adapted) value than it would without adaptation. This would simply add to the frequency of more moderate spatial frequency values reported, in a manner similar to the effect of guessing. However, the test Gabor was always presented at fixation, and fell therefore on a different part of the retina to the peripherally presented target Gabors, precisely to avoid perceptual interactions between the target and test stimuli.

We also ran similar simulations for the data obtained in Experiment 3 for position tracking, since the range of possible feature values here was also bounded by minima and maxima, though of course for position these boundaries were marked out in two dimensional space on the screen. These simulations included a data set where all 'responses' were simulated to occur at the mean recorded response location for that cage, as this is likely to be close to the most frequent location used as a guess. In addition, we used a simulated data set where all 'responses' were recorded as being near the most peripheral edge of cages, as this is a statistically unlikely location of a guess response. No differences in lags were observed in these simulations.

There is a further reason to conclude that the large lags obtained in Experiment 2 were not due to statistical patterns in trajectories of objects or to guessing strategies: in Experiment 7, we restricted the range of possible orientation values to between 0 and 45 degrees. In this control experiment, the range of orientation values was bounded by a minimum and a maximum value, making it similar to the case of spatial period in Experiment 2. This did not produce exaggerated lag patterns for orientation tracking which may have resulted from a combination of guessing and statistics of trajectories. Nor indeed do we observe lags of the magnitude seen in Experiment 2 for position tracking (Experiment 3), even though position values were similarly bounded, allowing for the possibility of guessing based on a mean feature value.